MIT Technology Review

VOL. 119 NO. 2 MARCH/APRIL 2016

10 Breakthrough Technologies

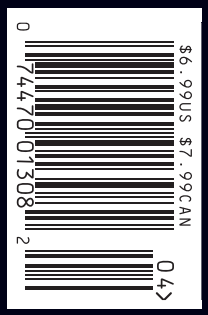

When will the smart, secure and seamless Internet of Things be a reality?

We started by connecting the phone to the Internet, now we're connecting the Internet to everything. By inventing technologies that connect your car, your home, and the cities in which we all live, we're accelerating a smarter, more seamless and intuitively synchronized world.

We are Qualcomm, and these are just a few of the ways we're bringing the future forward faster.

#WhyWait to join the discussion
Qualcomm.com/WhyWait

Why Wait™

QUALCOMM®

From the Editor

GUIDO VITTI

Lists, which are a staple of publishing because readers like them, also serve this purpose: they express what a publication thinks is important. Our annual list of 10 breakthrough technologies is a celebration of technological advances, but it's also a map of our emphases.

Sometimes we care about a technology because it offers new conveniences. "Tesla Autopilot," on page 62, describes how the carmaker sent its vehicles a software update that made autonomous driving a commercial reality. In "Conversational Interfaces" (see page 42), we report how Baidu's voice recognition and natural language software provide Chinese mobile-device users with practical speech interfaces. In "DNA App Store," on page 52, we write about a democratic approach to DNA sequencing in which consumers will share their genetic information to a common database and app developers will create products that use that information. Our profile of Slack (see page 60) accounts for the popularity of the workplace collaboration software by observing that the "ambient awareness" of an organization fostered by Slack's single stream of short, casual messages "creates the perception that keeping in touch with coworkers is effortless."

Sometimes we care about a breakthrough because new efficiencies significantly reduce the costs of a technology. The reusable rockets of SpaceX and Blue Origin (see page 46) will transform the economics of spaceflight, encouraging unimagined innovations. At SolarCity's gigafactory in Buffalo, New York (see page 54), the largest solar manufacturing facility in North America will make as many as 10,000 highly efficient solar panels every day. Efficiency matters because the panels themselves represent only 15 to 20 percent of the cost of installing a residential solar array. Much of the rest comes in what's known as balance-of-system costs. SolarCity's installation, says the company, will require one-third fewer panels to produce the same amount of electricity as conventional installations. Thus, according to Peter Rive, SolarCity's chief technology officer, the gigafactory "sets us up for a future where solar plus batteries is cheaper than fossil fuels."

Finally, sometimes we care about a technology just because it is an engineering marvel that solves a problem. Soon, radio signals will provide all the power a simple Internet device needs through a technique called passive Wi-Fi (see "Power from the Air," page 66). Robots will share data on how to pick up ordinary objects and become more dexterous (see "Robots That Teach Each Other," page 48). Immune engineering (see page 34) will provide effective ways to treat cancers. CRISPR-Cas9, which makes gene editing easy, will change what is meant by genetically modified crops, since the resulting plants contain no foreign DNA (see "Precise Gene Editing in Plants," page 40). A lab in China has used the technique to create a fungus-resistant wheat, researchers are using it on rice to boost yields, and a group in the U.K. has used it to tweak a gene in barley in an effort to make a drought-resistant variety.

All these technologies, which today seem so marvelous, will become commonplace. When the writer Ryan Bradley test-drove Tesla's Autopilot—or, rather, was test-driven—he "was amazed by how quickly I got used to it, how inevitable it began to feel," he says. "As a Tesla engineer told me … the thing that quickly becomes strange is driving a car without Autopilot." The future is like that: it becomes the present, with its own challenges and opportunities.

But write to me at jason.pontin@technologyreview.com and say what you think about this year's 10 technologies.

REPUBLIC OF TURKEY PRIME MINISTRY
INVESTMENT SUPPORT AND
PROMOTION AGENCY

Turkey: An Emerging Regional Startup Hub

As a technological innovation, the Internet was really still in its infancy at the close of the 20th century. Today, it is changing virtually every aspect of the way we live — and generating new business potential as well.

Technology giants such as Google, Facebook, and Amazon have spawned fledgling industries that abound with opportunities. Those opportunities, in turn, have captured the interest of angel investors, wealthy patrons willing to support the startup businesses envisioned by entrepreneurs who may have boundless ideas but limited means.

There are fewer than 1,000 angel investors in Turkey, according to the Turkish Prime Ministry's Treasury Department. That may seem small compared with the United States, where the Angel Capital Association says at least 300,000 angel investors are actively seeking investment prospects. However, to tap into the potential of its young and dynamic population, Turkey is taking new steps to increase the numbers of both entrepreneurs and angel investors while bolstering the country's nascent entrepreneurial ecosystem.

To that end, the new Turkish government has laid out an action plan for 2016 that addresses areas such as patent law, research and development support, and the generation of funds to support R&D and early-stage design work.

Turkey's entrepreneurial ecosystem is growing rapidly, with a variety of companies having recently secured investor support. The list includes: GittiGidiyor, an online auction/shopping platform; Yemeksepeti, the online food-ordering business; shopping portal Markafoni; Monitise, which develops technology for banking and other sectors; Vepa Grup, a cosmetics company; ReklamZ, an online advertising network; 4129Grey, an integrated digital-communication agency; Mynet, an Internet and media portal; Trendyol, an online fashion retailer; and AirTies, a wireless-network provider.

In addition, there are more platforms these days for entrepreneurs and others in Turkey who are looking to take their ideas to market. They can convene annually at Startup Istanbul, Startup Turkey, and the Webrazzi Summit to exchange ideas and experiences. They can also benefit from resources available via Turkey's 39 technoparks and 23 accelerators.

Turkey's evolving entrepreneurial ecosystem also includes nongovernmental organizations, business-angel networks, and venture capitalists. All these players are intensifying their efforts to nurture entrepreneurial growth.

So thanks to the ongoing efforts to nurture and support all aspects of entrepreneurship, Turkey is today poised to become a regional startup hub, ideally situated between Asia and Europe.

For more information, visit: invest.gov.tr

FINANCING ROUNDS – 2015 Q4, TURKEY

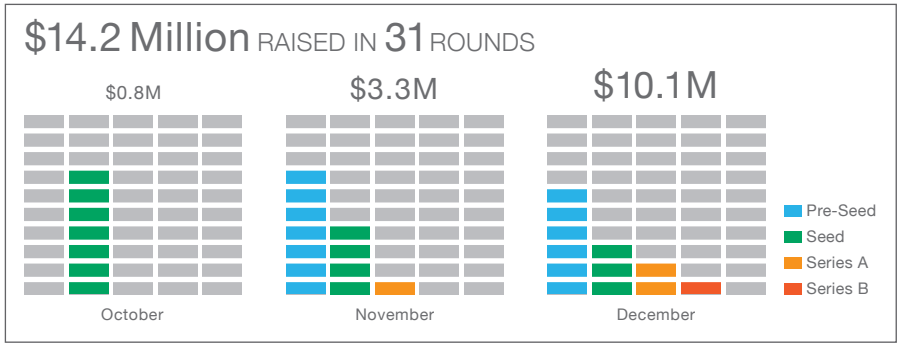

Contents

Front

MARCH/APRIL 2016

10 Breakthrough Technologies

New ways to cure disease, fly to space, and power our gadgets. New ways to communicate and to feed a growing population. These are the advances that will change the way you see the world, and the way the world works, in the coming years.

Back

ON THE COVER:

Design by Jessica Svendsen

Who will prevent downtime and equipment failure?

You and NI will. With an integrated platform that combines flexible, rugged hardware with intuitive software, NI helps organizations improve operational efficiency by providing systems that monitor and analyze rotating equipment. With advanced I/O, complex signal processing, and data analytics and visualization capabilities, NI puts you on the cutting edge of the Industrial Internet of Things and connects equipment, people, and technology like never before. See how at **ni.com/mcm**.

800 891 8841

NATIONAL INSTRUMENTS™

Technology Review.com/mustreads

Visit us online for daily news and analysis, Business Reports, magazine archives, and more.

COMPUTING

Next Big Test for AI: Making Sense of the World

A new database will gauge progress in artificial intelligence as computers try to grasp what's going on in scenes shown in photographs.

BIOMEDICINE
First Monkeys with Autism Created in China

BUSINESS
Big Auto Searches for Meaning Beyond Selling Cars

ENERGY
The Dubious Promise of Bioenergy Plus Carbon Capture

MOBILE
In Pursuit of an Affordable Tablet for the Blind

MATERIALS
Is 3-D Printing the Key to Cheap Carbon-Fiber Parts?

COMPUTING
Your Future Self-Driving Car Will Be Way More Hackable

FROM THE ARCHIVES

Can Computers Create Literature?

March 1998

Profiling the artificial intelligence that rocked the literary world.

TECHNOLOGYREVIEW.COM/AI1998

Never Charge Your Cell Phone Again

July 2004

Nanotech breakthroughs in solar cell technology made big promises.

TECHNOLOGYREVIEW.COM/CELL2004

Subscribers and MIT alumni can access our complete online article archive by activating their Insider accounts at technologyreview.com/activate. Not a subscriber or alum? Subscribe now at technologyreview.com/getinsider.

Has your company developed a technological discovery, invention or major break-through?

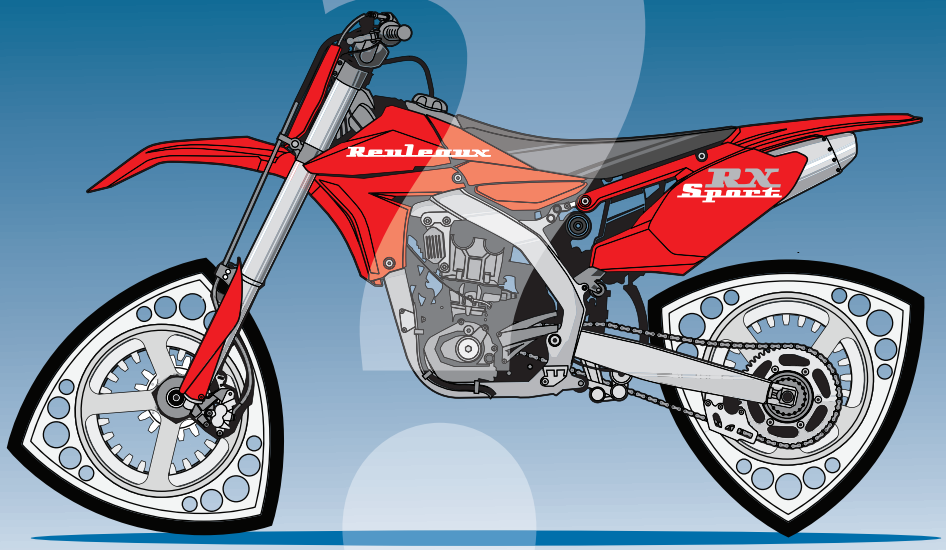

Why Reinvent the Wheel

when a tried and true solution has already stood the test of time?

Protect your brainstorm with the brain trust of **Allen, Dyer, Doppelt, Milbrath & Gilchrist!** Intellectual Property Law can be arcane and confusing to the uninitiated. If you don't have either an Intellectual Property Department or adequate expertise or resources, your I.P. legal issues could be best served by working with our nationally-recognized law firm, rather than attempting the work in-house. *All our attorneys are I.P. attorneys.*

Established in 1972, **ADDM&G** has protected the intrinsic rights of clients' original ideas and unique technologies through the application of patents, trademarks, copyrights, licensing, infringement, trade secrets and unfair competition protection, business litigation and counseling for 44 years. **ADDM&G** serves its clients' needs locally, statewide, across the nation and around the world. *Contact **ADDM&G** today for further information.*

ADDM&G

ALLEN, DYER, DOPPELT, MILBRATH & GILCHRIST

INTELLECTUAL PROPERTY ATTORNEYS

Orlando · 407.841.2330 | Jacksonville · 904.398.7000* | Miami · 305.374.8303
Tampa · 813.639.4222* | Winter Springs · 407.796.5064 | www.addmg.com

· Copyrights · Patents · Trademarks · Litigation ·

*Satellite Office

CUSTOMER SERVICE AND SUBSCRIPTION INQUIRIES
National: 800-877-5230

International: 903-636-1115

E-mail: technologyreview@pubservice.com

Web: www.technologyreview.com/customerservice

MIT Records: 617-253-8270 (alums only)

Reprints:
techreview@wrightsmedia.com
877-652-5295

Licensing and permissions:
licensing@technologyreview.com

Technology Review
One Main Street, 13th Floor
Cambridge, MA 02142
Tel: 617-475-8000

The mission of *MIT Technology Review* is to equip its audiences with the intelligence to understand a world shaped by technology.

Technology Review, Inc., is an independent nonprofit 501(c)(3) corporation wholly owned by MIT; the views expressed in our publications and at our events are not always shared by the Institute.

De Technologia non multum scimus. Scimus autem, quid nobis placeat.

MIT Technology Review

EmTech

EmTech conferences bring our award-winning journalism to life on stages around the globe, giving you access to the people and emerging technologies changing the world.

Asia | Brazil | Cambridge | China | Colombia | Ecuador | France

Hong Kong | India | Mexico | San Francisco | Spain

Attend the conference near you:
technologyreview.com/events

Feedback

E-mail letters@technologyreview.com

Write *MIT Technology Review*
One Main Street, 13th Floor
Cambridge, MA 02142

Please include your address, telephone number, and e-mail address.

Letters and comments may be edited for both clarity and length.

Five Most Popular Stories

MIT Technology Review
Volume 119, Number 1

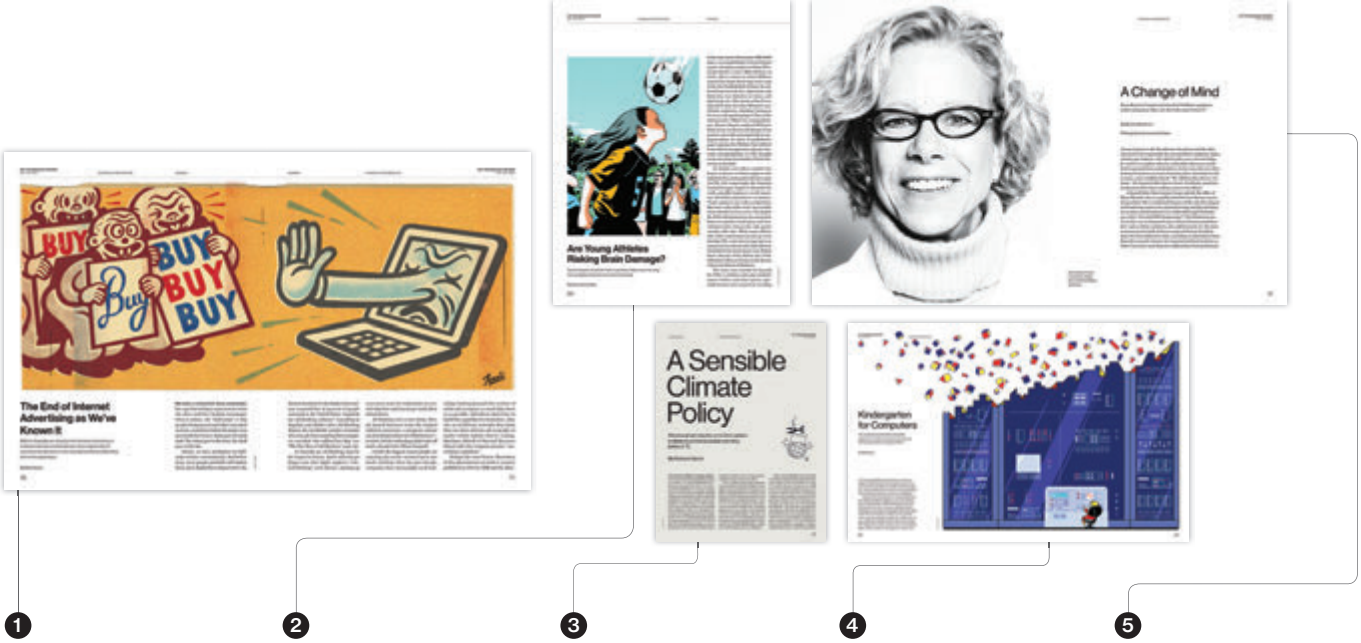

1 The End of Internet Advertising as We've Known It

The article is very pro ad-blocking, but we live in a society where we are given a lot of technology for free. There is an implicit collaboration going on between those that advertise and those that consume the "free" applications and information. Consider this before advocating for pulling the plug on the funding that has given so much benefit to the online community.

—**jrarchibald**

2 Are Young Athletes Risking Brain Damage?

Compare two similar sports: American football and rugby. The basic rules are very similar. One experiences order-of-magnitude fewer head trauma injuries at every competitive level: rugby. What is the crucial difference? No protective gear to speak of in rugby. No pads, no helmets. Hence, the solution is almost trivial: remove protective gear. Players will instinctively adjust their behavior to the acceptable risk levels. This won't eliminate all injuries, but it will dramatically reduce the number and severity of injuries.

—**dusanmal**

3 This Climate Policy Could Save the Planet

The problem with Martin's approach is that it requires politicians to take aggressive actions—such as drastically limiting emissions in poor countries—that may be unrealistic and even harmful.

—**hemanth chowdary**

@hemanth chowdary
In rich countries, "drastically limiting emissions" should not be so damaging, especially when there is so much waste, including food—and we shouldn't forget that agriculture is the primary source of methane emissions. —**lacramioara.astefanoaei**

4 Kindergarten for Computers

We seem to be going up a blind alley here. The human brain evolved over millennia, building error on error, finding ways around architecture limitations and often rigging the system to get the Darwinistic result of survival. It never evolved to be the best thinker for the complex modern world.

—**datascienceox**

The development of individual human intelligence relies heavily on interaction with other intelligence. We ought to have all those AIs "talk" to each other more often. —**yuanbo**

5 A Change of Mind

As a mother of a seven-year-old girl with Down syndrome, lacking education in this field, I would like to share my observation: sometimes my daughter Elma surprises us with her observation, conclusions, and wits. It seems like a ray of light exposes her cognitive abilities. Then it goes away.

—**NadzorMama**

Climate Change Head Scratcher

I find it hard to understand why climate change mitigation should be so difficult and failing so badly, and why solutions elude such a well-educated audience of scientists and engineers.

For the near and intermediate terms we have two choices: either drastically reduce energy consumption and "first world" standards of living (politically impossible) or rapidly expand carbon-free energy sources, such as hydroelectric and nuclear, that are already in full-scale use, with known risks and economics.

Ken Caldeira ("Stop Emissions!" January/February 2016) doesn't even mention that there are developable, infinitely renewable hydroelectric resources worldwide, including in North America.

Nuclear has its problems, but Three Mile Island did, in fact, validate the basic design principles of redundant control and reactor containment. Ignoring these principles caused Chernobyl. Fukushima was caused by a blatant disregard for common-sense design criteria. With these lessons learned, and with proposed new reactor designs, we have the basis for a renewed, safe nuclear power industry.

Caldeira doesn't even mention that there are developable, infinitely renewable hydroelectric resources worldwide.

For long-term success, we must continue basic and applied research on solar, wind, nuclear fusion, and still-to-be-discovered advances. We must avoid the panicked waste of huge resources. In particular, engineers should understand the fallacy of applying the dot-com venture capital model to large-scale chemical and mechanical manufacturing facility development. We need well-defined products of scientific research followed by pilot-scale demonstration, development, and process refinement followed by full-scale deployment.

To accomplish all this, the first thing we need is support from the scientific and technical communities, followed by public policy changes to reposition the incentives and disincentives. Why is this so difficult?

—David Korenstein, Wayne, Pennsylvania

Views

Greg Lennon

Julie Morgenstern

Ashutosh Saxena

BIOMEDICINE

Own Your DNA

Having your genome sequenced doesn't always mean you have full access to the data.

We've long been promised the chance to have our genomes sequenced at a reasonable price, and now we're finally at the point where companies worldwide are launching large-scale sequencing services at prices likely to attract consumers.

But once you've paid to have your genome sequenced, will you have full, downloadable access to it?

The general principle would seem to be simple: you have a basic right to the data derived from your own body. Many genome pioneers assumed that principle was obvious when they sought and received public support for the Human Genome Project. Since then, entities ranging from the U.S. Food and Drug Administration to the White House have reinforced the idea that people should have access to their own genomic data.

And yet not everybody seems ready to line up behind it. There are now companies as well as large-scale government projects that either are blocking total downloadable access or appear to be wavering on access in the future. They range from well-known government projects compiling thousands of genomes to smaller companies that are promising to set up an app ecosystem where customers get bits and pieces of their genomes upon request (see "DNA App Store," page 52).

I think this is a questionable path to take, and here are a few reasons why:

For one, errors are bound to occur with sequencing, and letting customers have full access to their data allows them to compare data generated at different times or by other means or other labs.

Second, allowing the customer to download the data can help sequencing companies avoid liability, since they can't be said to have prevented a user from learning about an important variation.

Third, having access to your data means you can share it—with a family member, a health-care provider, a genetic counselor, a citizen science group, a patient group, or a research group. Pooling data in this way is hugely important in large-scale genomic studies.

Fourth, you can transport it as you wish. What we know today about our genomes pales in comparison with what we'll know in the future. As we discover new associations for variants, our genomes become more meaningful. There will never be only one interpretation of a genome, so it's important to be able to download and digitally transport it if you feel you need a second or third opinion.

We need to assert the right to our DNA data now, before any bad precedent is set in the other direction. That data will become even more valuable and important as genome sequencing becomes ever more widespread.

Greg Lennon, who has over 25 years of experience as a human genome scientist, is the cofounder of SNPedia and Promethease.

COMPUTING

E-Mail's Not the Issue

Why a new messaging platform can't solve our productivity problems.

I'm a productivity consultant, which means I work every day with people to organize their time, and to make the best use of whatever information and energy they have. I've been doing this for 25 years, so I was around when e-mail was the great savior that would make us all more organized and productive. Instead, it's crushing our ability to get anything done. Clients tell me: "I don't have any

ANDY FRIEDMAN

time to think." "It's hard to turn work off at night." "I can't focus."

Why is e-mail so bad? It's like an overstuffed toy box full of miscellaneous stuff—missives from your boss, requests from clients, FYIs, news feeds, social-media alerts, junk. We spend more time sifting through it all than getting anything done.

It's led to an instant-response culture, with pressure to constantly check messages. It's no longer okay to reply a few hours from now. You're expected to interrupt your work and respond *now*. And then there's the fear-of-missing-out phenomenon. Even if you're "allowed" to turn e-mail off, you don't, because of competitive insecurity: if your boss puts something out there and your colleagues respond first, you feel you've lost out.

So along comes an app like Slack, which offers to combat some of e-mail's perils (see "Slack," page 60). By separating internal company communications from the general "toy box," it offers something less miscellaneous, more focused, more collaborative. Slack's "channel" feature consolidates information by topic, offering a snapshot of a project's status—in theory much better than rummaging through e-mail threads to find what you need.

But like all technologies, Slack is a tool whose success depends on the skill of the people using it. How many iterations of "Thanks!" and "Got it!" do you have to weed through to get to the substantive info? Without clear protocols, Slack channels can get just as messy as e-mail.

Companies need to actively shift away from the instant-response culture. How? Establish protocols for reasonable turnaround times. Make it highly uncool to send things at the last minute. Create quiet spaces and hours. Train people to disconnect and think—to strategize, analyze, problem-solve, and recharge. People can't perform at their best on high-level tasks by squeezing those tasks into seven-minute windows between checking messages—and it doesn't really matter if those messages are coming over Slack, e-mail, or something else altogether.

Julie Morgenstern is the author of, among other books, Never Check E-Mail in the Morning and Other Unexpected Strategies for Making Your Life Work.

ROBOTICS

Wikipedia for Robots

People have learned to pool their knowledge. We need to help machines do the same.

Humans have gained a lot of value by organizing all their knowledge and making it widely accessible—in textbooks, libraries, Wikipedia, and YouTube, to name a few examples. These pools of knowledge aren't valuable just for grand scientific ventures but also for the trivial stuff of everyday human lives: you can easily find thousands of YouTube videos that will teach you how to cook an omelet.

We now live in a world where robots are helping humans in their daily lives, and just like humans, robots need to learn new skills in order to do their jobs successfully. And we shouldn't expect a robot to learn on its own from scratch, any more than we'd expect a human to do so—imagine a child growing up with no access to textbooks, libraries, or the Internet.

However, the organized collections of knowledge that work for humans aren't so great for robots. A robot wouldn't get much useful information if it queried a search engine for how to "bring sweet tea from the kitchen." Robots require something different—access to finer details for planning, control, and natural language understanding. When asked to bring sweet tea, the robot would need access to the knowledge for interpreting the language symbols ("tea") in terms of physical entities ("a particular container having sweet tea"), the spatial knowledge that sweet tea can be either on a table or in a fridge, and the knowledge for inferring how to grasp and manipulate objects. It's possible to manually script a demo for one particular situation, but handling this across different tasks and in different environments is still an open problem.

In 2014, I started a project called RoboBrain at Cornell University along with PhD students Ashesh Jain and Ozan Sener. We now have collaborators at Stanford and Brown. What we're working on is a way of sharing information that allows robots to gather whatever knowledge they need for a task (see "Robots That Teach Each Other," page 48). If one robot learns, then the knowledge is propagated to all the robots. RoboBrain achieves this by gathering the knowledge from a variety of sources. The system stores multiple kinds of information, including symbols, natural language, visual or shape features, haptic properties, and motions.

This approach represents a huge shift in thinking. Historically, research groups working with robots have trained their robots in isolation. Yes, we often share ideas through publications and software that can be used by another research group, but what one robot might learn hasn't been accessible to another researcher's robot. To add to the problem, research groups have been working on different problems—one might have focused on the computer vision problem of identifying a cup, while another worked on the language problem of what is a "cup," while a third tackled how to grasp a cup.

That's the kind of approach we need to get past. A cup is one object, not three. And a robot, just like a person, needs to be able to have all the knowledge it needs in one place.

Ashutosh Saxena is the director of the RoboBrain project and the founder and CEO of the startup Brain of Things.

The All New
technologyreview.com

New Design
The best reading experience on all your devices.

New Features
The Daily – Every day, new must-read technology stories in one convenient location.

More of the technology news and analysis you need.

New Access
The complete archive of online stories since 1997.

Experience it today.

MIT Technology Review

technologyreview.com

Upfront

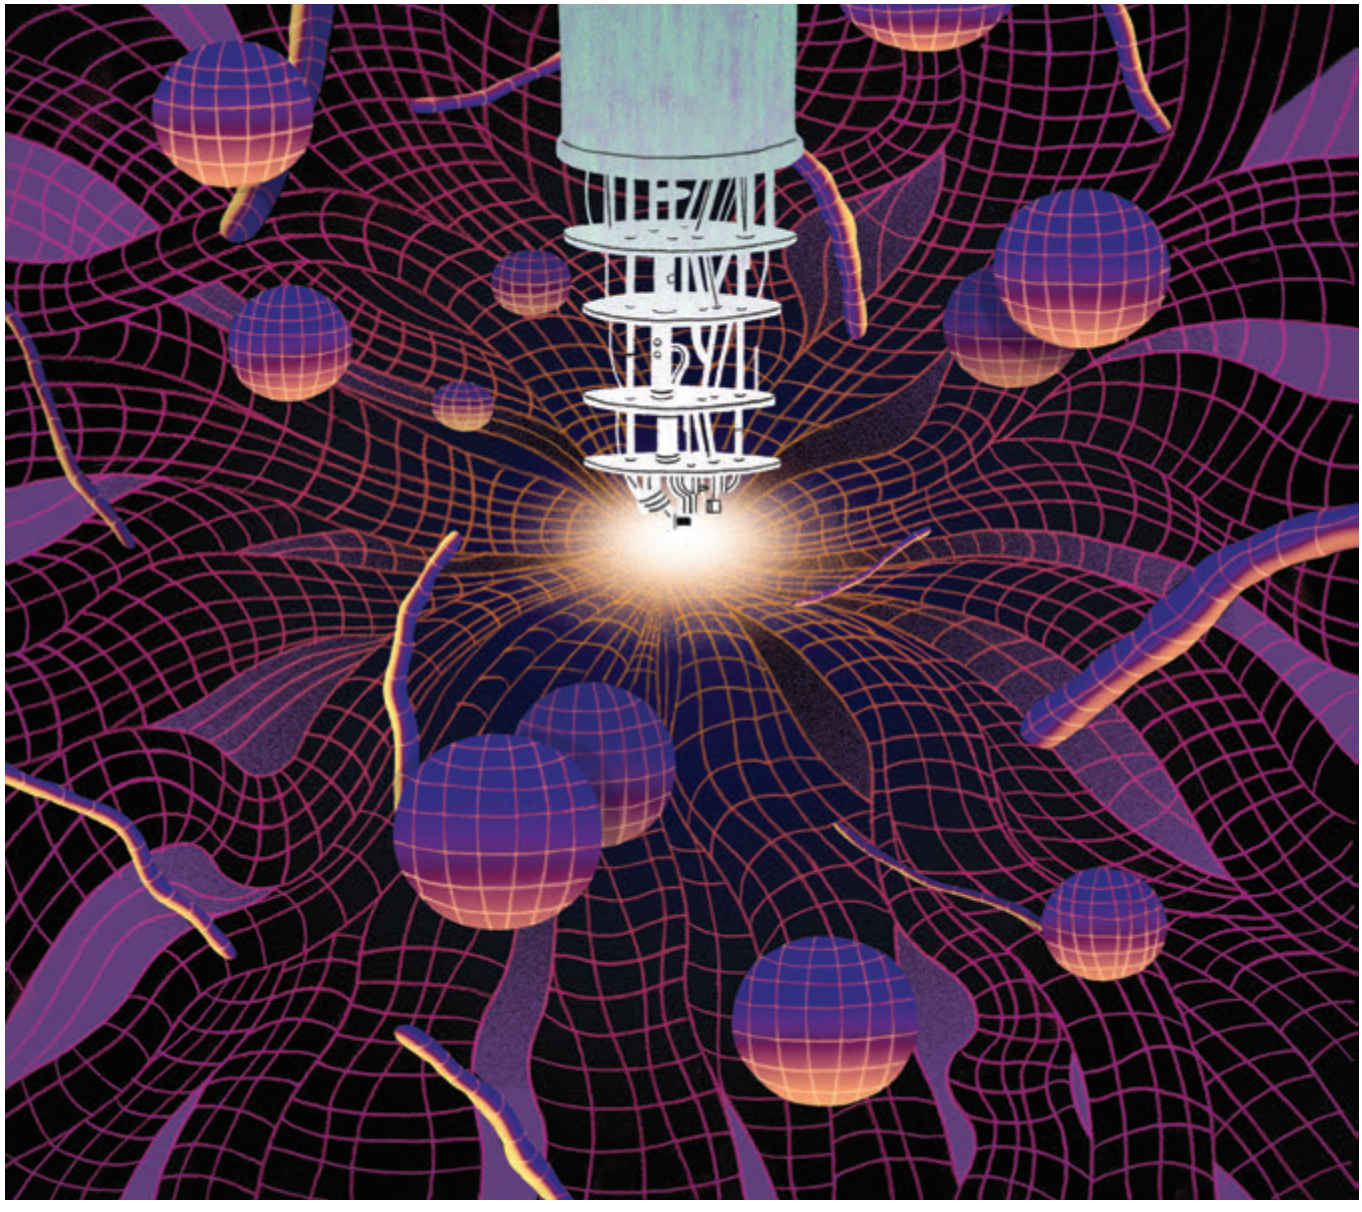

DADU SHIN

Google's Quantum Dream Machine

Physicist John Martinis could deliver one of the holy grails of computing to Google—a machine that dramatically speeds up today's applications and makes new ones possible.

John Martinis used the arm of his read-ing glasses to indicate the spot where he intends to demonstrate an almost unimaginably powerful new form of computer in a few years. It is a cylindrical socket an inch and a half across, at the bottom of a torso-sized stack of plates, blocks, and wires of brass, copper, and

Upfront

gold. The day after I met with him last fall, he loaded the socket with an experimental superconducting chip and cooled the apparatus to a hundredth of a degree Celsius above absolute zero. To celebrate that first day of testing the machine, Martinis threw what he called "a little party" at a brew pub with colleagues from his Google lab in Santa Barbara, California.

No doubt a much bigger celebration will take place if Martinis and his group can actually create the wonder computer they seek. Because it would harness the strange properties of quantum physics that arise in extreme conditions like those on the ultracold chip, the new computer would let a Google coder run calculations during a coffee break that would take a supercomputer of today millions of years. The software that Google has developed on ordinary computers to drive cars or answer questions could become vastly more intelligent. And earlier-stage ideas bubbling up at Google and its parent company, such as robots that can serve as emergency responders or software that can converse at a human level, might become real.

The theoretical underpinnings of quantum computing are well established. And physicists can build the basic units, known as qubits, out of which a quantum computer would be made. They can even operate qubits together in small groups. But they have not made a fully working, practical quantum computer.

Martinis's research group at the University of California, Santa Barbara, has demonstrated some of the most reliable qubits around and gotten them running some of the code a quantum computer would need to function. He was hired by Google in June 2014. With his new Google lab up and running, Martinis guesses that he can demonstrate a small but useful quantum computer in two or three years.

Martinis and his team have to be adept at many things because qubits are fickle.

"We often say to each other that we're in the process of giving birth to the quantum computer industry," he says.

As recently as December the prospect of a quantum computer doing anything useful within a few years seemed remote. Researchers in government, academic, and corporate labs were far from combining enough qubits to make even a simple proof-of-principle machine. A well-funded Canadian startup called D-Wave Systems sold a few of what it called "the world's first commercial quantum computers" but spent years failing to convince experts that the machines actually were doing what they should.

Then, last December, NASA summoned journalists to building N-258 at its Ames Research Center in Mountain View, California, which hosts a D-Wave computer bought by Google. There Hartmut Neven, who leads the Quantum Artificial Intelligence lab Google established to experiment with its D-Wave machine, unveiled the first real evidence that it can offer the power that proponents of quantum computing have promised. In a carefully designed test, the superconducting chip inside D-Wave's computer—known as a quantum annealer—performed 100 million times faster than a conventional processor.

However, this kind of advantage needs to be available in practical computing tasks, not just contrived tests. That's where Martinis comes in. Neven doesn't think D-Wave can get a version of its quantum annealer ready to serve Google's engineers quickly enough, so he hired Martinis to do it. "It became clear that we can't just wait," Neven says. "There's a list of shortcomings that need to be overcome in order to arrive at a real technology."

Google will be competing not only with whatever improvements D-Wave can make but also with Microsoft and IBM, which have substantial quantum computing projects of their own. But those companies are focused on designs

TO MARKET

VirZoom

Virtual-reality bike

COMPANY:
VirZoom

PRICE:
$250

AVAILABILITY:
First half of 2016

COURTESY OF VIRZOOM

A new stationary bike from Boston startup VirZoom requires an unusual accessory while you're pedaling: a virtual-reality headset to add some adventure to your spinning. The bike looks like a traditional folding one except for the buttons and triggers scattered across its two handles. VirZoom cofounder and CEO Eric Janszen recently challenged me to a race-car race inside the headset. I leaned to hug corners on tight turns and pedaled faster to speed up my car. When I approached rougher ground, I was forced to pedal harder to keep up the same pace. It's compatible with three forthcoming virtual-reality headsets—the Oculus Rift, Sony PlayStation VR, and HTC Vive. *—Signe Brewster*

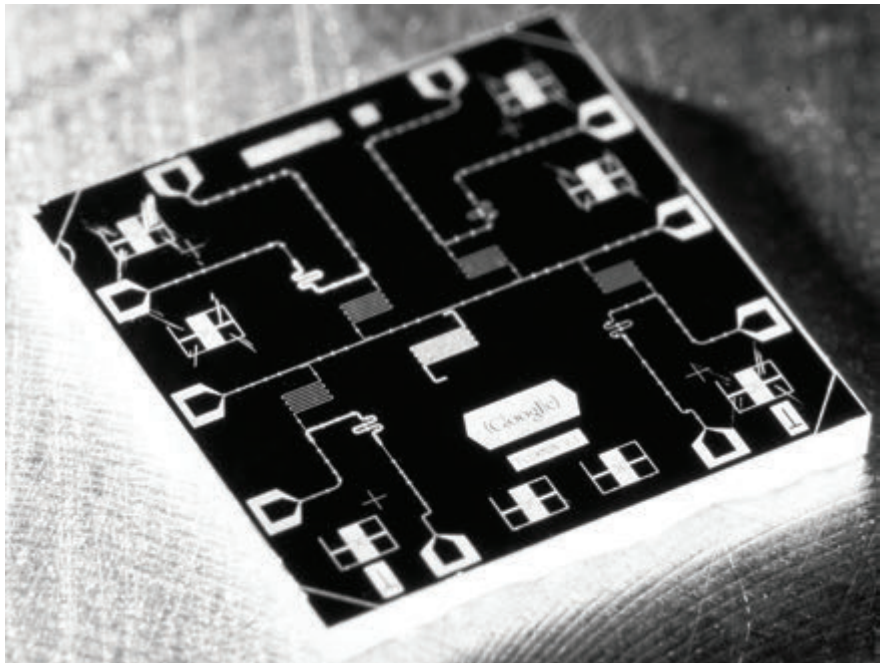

An experimental chip is cooled to near absolute zero in order to generate quantum effects.

COURTESY OFGOOGLE AND SPENCER BRUTTIG

much further from becoming practically useful. Indeed, a rough internal time line for Google's project estimates that Martinis's group can make a quantum annealer with 100 qubits as soon as 2017. D-Wave's latest chip already has 1,097 qubits, but Neven says a high-quality chip with fewer qubits will probably be useful for some tasks nonetheless. A quantum annealer can run only one particular algorithm, but it happens to be one well suited to the areas Google most cares about. The applications that could particularly benefit include pattern recognition and machine learning, says William Oliver, a senior staff member at MIT Lincoln Laboratory who has studied the potential of quantum computing.

Martinis and his team have to be adept at so many things because qubits are fickle. They can be made in various ways—Martinis uses aluminum loops chilled until they become superconducting—but all represent data by means of delicate quantum states that are easily distorted or destroyed by heat and electromagnetic noise, potentially ruining a calculation.

The difficulty of creating qubits that are stable enough is the reason we don't have quantum computers yet. But Martinis thinks he's nearly there. The coherence time of his qubits, or the length of time they can maintain a superposition, is tens of microseconds—about 10,000 times the figure for those on D-Wave's chip. Martinis's confidence in his team's hardware even has him thinking he can build Google an alternative to a quantum annealer that would be even more powerful. A universal quantum computer, as it would be called, could be programmed to take on any kind of problem, not just one kind of math. The theory behind that approach is actually better understood than the one for annealers, in part because most of the time and money in quantum computing research have been devoted to universal quantum computing. A year ago Martinis and his team became the first to demonstrate qubits that crossed a crucial reliability threshold for a universal quantum computer. "We demonstrated the technology to a point where I knew we could scale up," says Martinis. "This was for real." He now believes he can show off a complete universal quantum computer with about 100 qubits around the same time he delivers Google's new quantum annealer, in about two years. That would be a major milestone in computer science.

When Martinis explains why his technology is needed at Google, he doesn't spare the feelings of the people working on AI. "Machine-learning algorithms are really kind of stupid," he says. "They need so many examples to learn."

Indeed, machine learning is pathetic next to the way humans or animals pick up new skills or knowledge. Teaching a piece of software new tricks, such as how to recognize cars and cats in photos, generally requires thousands or millions of carefully curated and labeled examples. Although a technique called deep learning has recently produced striking advances in the accuracy with which software can learn to interpret images and speech, more complex faculties like understanding the nuances of language remain out of machines' reach.

Figuring out how Martinis's chips can make Google's software less stupid falls to Neven. He thinks that the prodigious power of qubits will narrow the gap between machine learning and biological learning—and remake the field of artificial intelligence. "Machine learning will be transformed into quantum learning," he says. That could mean software that can learn from messier data, or from less data, or even without explicit instruction.

—Tom Simonite

Upfront

Have Global CO_2 Emissions Peaked?

New estimates suggest that worldwide emissions dropped slightly from 2014 to 2015. If so, that would be the first such decline that did not coincide with an economic downturn. But it's not yet clear that we can expect consistent economic growth without emissions growth.

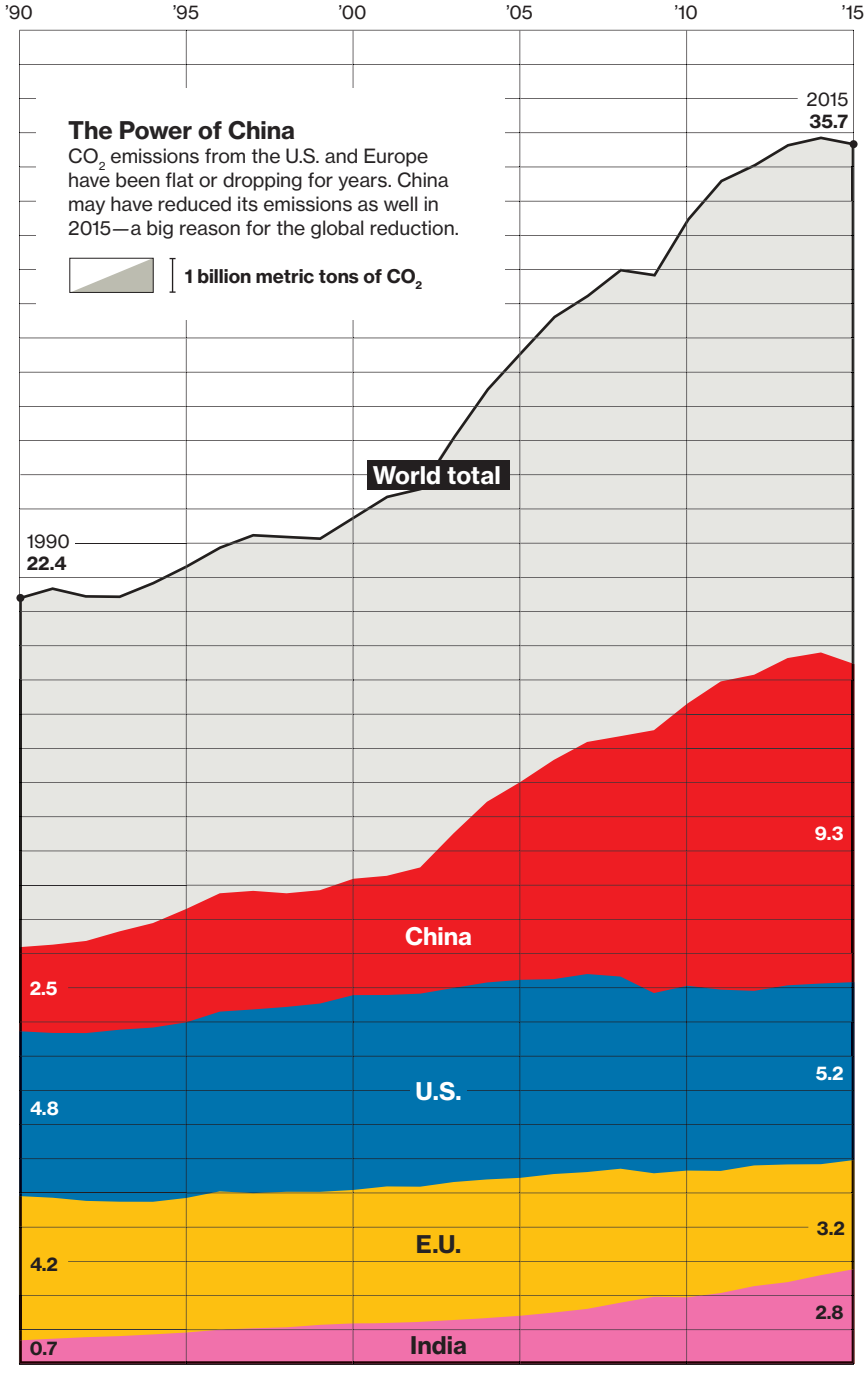

Breaking the Link

Emissions in the United States and Europe have fallen or leveled off even as per capita gross domestic product has increased.

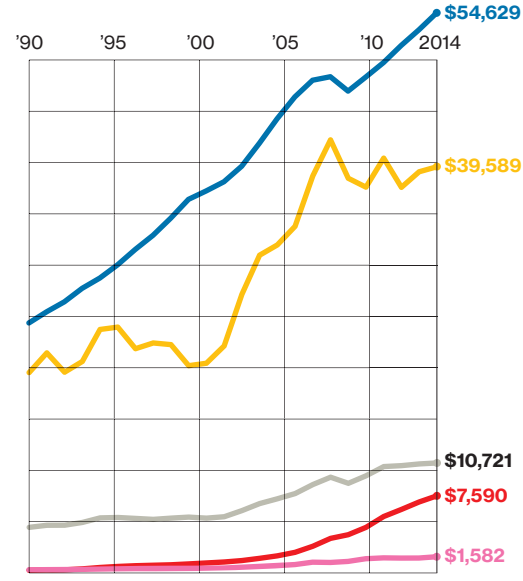

A Developing Situation

The emissions trajectory depends greatly on poor countries, where Harvard researchers expect the fastest annual growth through 2024.

1. India	7.0%
2. Uganda	6.0%
3. Kenya	6.0%
4. Tanzania	6.0%
5. Egypt	5.8%
6. Madagascar	5.8%
7. Senegal	5.8%
8. Philippines	5.7%
9. Malawi	5.7%
10. Zambia	5.6%
11. Guatemala	5.2%
12. Pakistan	5.1%
13. Zimbabwe	4.9%
14. Malaysia	4.9%
15. Indonesia	4.8%

ILLUSTRATION BY LUKE SHUMAN; DATA FROM THE GLOBAL CARBON PROJECT, THE WORLD BANK, NATURE CLIMATE CHANGE, AND THE HARVARD CENTER FOR INTERNATIONAL DEVELOPMENT

MITSloan
CIO Symposium

WEDNESDAY MAY 18, 2016

FULL DAY EVENT | CAMBRIDGE, MA

THRIVING IN THE DIGITAL ECONOMY

THE 13TH ANNUAL MIT SLOAN CIO SYMPOSIUM

The MIT Sloan CIO Symposium combines the academic thought leadership of MIT with the in-the-trenches experience of global CIOs and industry experts. It is the premier international conference for CIOs and business leaders to look beyond day-to-day issues and explore enterprise innovations in technology and business practices.

THIS YEAR'S TOPICS INCLUDE: The changing nature of work, the workplace and innovation; Big Data 2.0 and data strategy; platform strategies; IoT; cybersecurity; and blockchain.

This MIT all-star lineup consists of scholars who are affiliated with the MIT Initiative on the Digital Economy (IDE) and the MIT Sloan Center for Information Systems Research (CISR):

- Erik Brynjolfsson
- Kristine Dery
- Nils O. Fonstad
- Andrew P. McAfee
- Alex Pentland
- Marshall Van Alstyne
- Stephanie L. Woerner
- Peter Weill
- George Westerman
- Barbara H. Wixom

MITSloan MANAGEMENT CISR

MIT INITIATIVE ON THE DIGITAL ECONOMY

REGISTER NOW:
www.mitcio.com

Upfront

The Chimera Contention

A radical new approach to generating human organs is to grow them inside pigs or sheep.

Braving a funding ban put in place by America's top health agency, some U.S. research centers are moving ahead with attempts to grow human tissue inside pigs and sheep with the goal of creating hearts, livers, or other organs needed for transplants. The effort to incubate organs in farm animals is ethically charged because it involves adding human cells to animal embryos in ways that could blur the line between species.

Last September, in a reversal of earlier policy, the National Institutes of Health announced it would not support studies involving such "human-animal chimeras" until it had reviewed the scientific and social implications more closely. The agency, in a statement, said it was worried about the chance that animals' "cognitive state" could be altered if they ended up with human brain cells. The NIH action was triggered after it learned that scientists had begun such experiments with support from other funding sources, including California's state stem-cell agency. The human-animal mixtures are being created by injecting human stem cells into days-old animal embryos and then gestating these in female livestock.

The extent of the research was disclosed in part during presentations made at the NIH's Maryland campus in November at the agency's request. One researcher, Juan Carlos Izpisua Belmonte of the Salk Institute, showed unpublished data on more than a dozen pig embryos containing human cells. Another, from the University of Minnesota, provided photographs of a 62-day-old pig fetus in which the addition of human cells appeared to have reversed a congenital eye defect.

Desperately ill people on organ waiting lists might someday order a chimera.

The experiments rely on a cutting-edge fusion of technologies, including recent breakthroughs in stem-cell biology and gene-editing techniques. By modifying genes, scientists can now easily change the DNA in pig or sheep embryos so that they are genetically incapable of forming a specific tissue. Then they add stem cells from a person and hope that the human cells will take over the job of forming the missing organ, which could then be harvested from the animal for use in a transplant operation.

"We can make an animal without a heart. We have engineered pigs that lack skeletal muscles and blood vessels," says Daniel Garry, a cardiologist who leads a chimera project at the University of Minnesota. While such pigs aren't viable, they can develop properly if a few human stem cells are added.

Others worry that the animals might turn out to be a little too human for comfort. "We are not near the island of Dr. Moreau, but science moves fast," NIH ethicist David Resnik said during the agency's November meeting. "The specter of an intelligent mouse stuck in a laboratory somewhere screaming 'I want to get out' would be very troubling to people."

The chance of an animal gaining human consciousness is probably slim; their brains are just too different, and much smaller. Even so, as a precaution, researchers working with farm-animal

PING ZHU

chimeras haven't yet permitted any to be born, but instead are collecting fetuses in order to gather preliminary information about how great the contribution of human cells is to the animals' bodies.

Hiromitsu Nakauchi, a stem-cell biologist at Stanford University, began trying to make human-sheep chimeras last year. He says that so far the contribution by human cells to the animals' bodies appears to be small. "If the extent of human cells is 0.5 percent, it's very unlikely to get thinking pigs or standing sheep," he says. "But if it's large, like 40 percent, then we'd have to do something about that."

The process for creating chimeras, called "embryo complementation," involves placing human cells into an animal embryo at the very earliest stage, when it is a sphere of just a dozen cells in a laboratory dish. The human cells can multiply, specialize, and potentially contribute to any part of the animal's body as it develops. In 2010, while working in Japan, Nakauchi used the embryo complementation method to show he could generate mice with a pancreas made entirely of rat cells. Although he was a star scientist, Japanese regulators were slow to approve his idea for chimeras, and by 2013 Nakauchi decided to move to the U.S., where no federal law restricts their creation. Stanford was able to recruit him with the help of a $6 million grant from the California Institute of Regenerative Medicine.

The type of human cells Nakauchi's team adds to animal embryos in the Stanford lab, called iPS cells, are made from skin or blood cells chemically reprogrammed into more versatile stem cells. Nakauchi says that most of the iPS cells his team has been placing into animal embryos are made from his own blood, since recruiting volunteers involves too much paperwork. If his iPS cells develop inside an animal, the resulting tissue will actually be his, a kind of perfectly matched replacement part. Desperately ill people on organ waiting lists might someday order a chimera and wait less than a year for a custom organ to be ready. "I really don't see much risk to society," Nakauchi says.

Before that can happen, scientists will have to prove that human cells can really contribute effectively to the bodies of farm animals. To find out, researchers in 2014 decided to begin impregnating farm animals with human-animal embryos, says Pablo Ross, a veterinarian and developmental biologist at the University of California, Davis. Ross says he has transferred about six sets of pig-human embryos into sows in collaboration with the Salk Institute and worked with Nakauchi to establish another eight or 10 pregnancies involving sheep-human embryos.

Hiromitsu Nakauchi

These early efforts aren't intended to make organs, says Ross, but more "to determine the ideal conditions for generating human-animal chimeras." The embryonic animals grow to half an inch long, just enough development to see if human cells are contributing to the organs. "My view is that the contribution of human cells is going to be minimal, maybe 3 or 5 percent," says Ross. "But what if the embryo that develops is mostly human? It's something that we don't expect, but no one has done this experiment, so we can't rule it out." *—Antonio Regalado*

QUOTED

"AI technology can solve problems that scale to the whole planet."

—Facebook CTO Mike Schroepfer on how artificial intelligence can help address global issues like poverty and climate change.

"It's amazing what they do to their brains."

—Ann McKee, professor of neurology and pathology at the Boston University School of Medicine, who is studying the brains of dead football players to learn more about head trauma and neurodegenerative disease.

"We can't wait an extra decade to bring new nuclear plants online."

—Terrestrial Energy CEO Simon Irish on the need to overcome regulatory barriers for advanced nuclear technology.

BY THE NUMBERS

$599

Price of the Oculus Rift VR headset without the additional cost of the computer it needs to work.

300 gigabits per second per square millimeter

Rate at which a prototype microprocessor using optical connections instead of electrical wires transmits data.

90%

Degree to which a gyroscope-equipped glove for Parkinson's patients reduced hand tremors in tests.

0.84 seconds

Time it takes trained drivers to regain control of Google's autonomous car from its computer, on average.

Upfront

Smart Bandages

Researchers have developed a new kind of wound dressing that could serve as an early-detection system for infections.

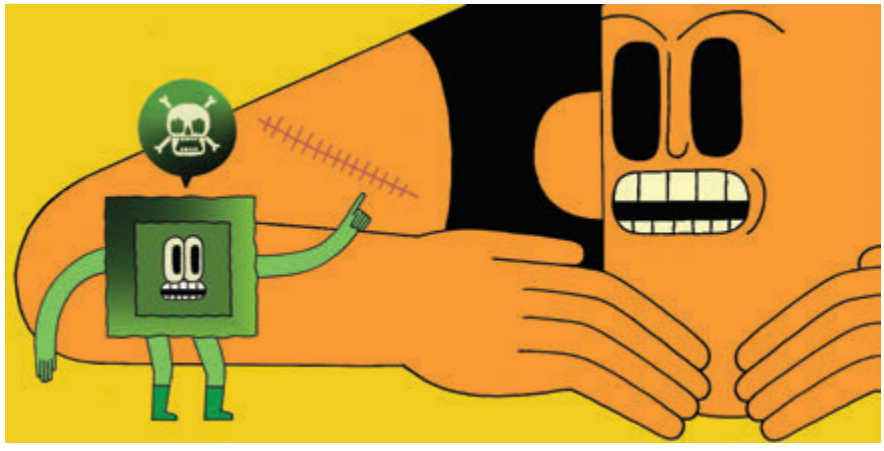

Bacterial infection is a fairly common and potentially dangerous complication of wound healing, but a new "intelligent" dressing that turns green to signal the onset of an infection could serve as a valuable early-detection system.

Researchers in the United Kingdom recently unveiled a prototype of the bandage, which contains a gel-like material infused with tiny capsules that release nontoxic fluorescent dye in response to contact with populations of bacteria that commonly cause wound infections. Led by Toby Jenkins, a professor of biophysical chemistry at the University of Bath, the inventors of the new bandage, which has not yet been tested in humans, say it could be used to alert health-care professionals to an infection early enough to prevent the patient from getting sick. In some cases it might even help avoid the need for antibiotics.

Jenkins's group is collaborating with clinical researchers from a pediatric burn center at the University of Bristol. Clinicians tend to overprescribe antibiotics for burn wounds, particularly in children, because they are so concerned about infection. An infection-detecting bandage could reassure parents and doctors when a wound is in fact not infected. It would also be useful for monitoring surgical wounds as well as those that result from traumatic injury, says Jenkins.

All wounds get colonized by bacteria, often including pathogenic species, but small populations are generally not harmful. In some cases, though, a population of harmful bacteria grows too big for the immune system to handle, and clinical intervention is needed to clear it. "We believe that this transition normally happens several hours, if not longer, before any clinical symptoms become evident," says Jenkins. Earlier detection might give doctors time to head off the infection even before such symptoms arise.

Jenkins says the transition is "almost certainly" associated with the formation of a so-called biofilm, a layer of microbes that work together and secrete a slimy substance to defend the colony against the immune system. At a high enough population density, the bacterial film switches on the production of toxins, which puncture the capsules as they would cells in the body, releasing the dye, which fluoresces when it is diluted by the surrounding gel.

Though the clinical utility of the bandage has yet to be proven, Jenkins says the technology could be ready for clinical testing by 2018. *—Mike Orcutt*

TO MARKET

The 184

Passenger drone

COMPANY:
Ehang

PRICE:
$200,000–$300,000

AVAILABILITY:
2016

Ready to commute via drone? The 184, built by Chinese consumer drone maker Ehang, is a 440-pound quadcopter with an enclosed seating area for human passengers. For about the price of a small helicopter, the 184 can carry a person up to 10 miles, or up to 23 minutes, at speeds around 60 miles per hour. It's capable of flying at heights of up to 2.15 miles, though drone regulations would likely keep it much lower than that. Made of lightweight carbon fiber, the drone takes off and lands vertically and is battery-powered. A set of flight-control algorithms oversee the speed of the drone's huge rotors. Navigation is as simple as putting an address in Google Maps. *—Signe Brewster*

JOHN MALTA; COURTESY OF EHANG

PROFESSIONAL EDUCATION

MASSACHUSETTS INSTITUTE OF TECHNOLOGY

EXPAND YOUR KNOWLEDGE WITH MIT PROFESSIONAL EDUCATION PROGRAMS

Looking for a way to boost your knowledge and advance your career? MIT Professional Education gives technical professionals an added edge by providing crucial and timely knowledge in specialized science and engineering fields through an array of program offerings. Taught by world-renowned faculty from across the Institute, these programs provide a gateway to MIT expertise and give participants the opportunity to further their careers, elevate organizational performance and help make a difference in the world.

MIT Professional Education includes intensive one to five day programs, self-paced digital programs, full-semester on-campus MIT courses, custom programs at your company location and even short courses at international locations.

▶ DIGITAL PROGRAMS

Bring MIT to You

Digital Programs allow busy professionals globally to take courses taught by MIT faculty, from any location and at their own convenience. These online programs focus on important industry topics relevant to a variety of professionals around the world. Register now and save 10% with code TR10. Upcoming offerings include:

- **Internet of Things: Roadmap to a Connected World** April 12, 2016–May 24, 2016
- **Entrepreneurial Negotiations: The MIT Way** April 26, 2016–June 7, 2016

▶ SHORT PROGRAMS

Come to MIT for a Week

Short Programs offer more than 50 intensive programs addressing new and evolving technologies, challenges and breakthrough solutions. Courses combine learning from MIT's groundbreaking research with real-world application perspectives from industry, government and academia. Register by April 15 to save 10% with code TR10. Earn CEUs and a Certificate of Completion in a variety of topic areas:

- Biotechnology / Pharmaceutical
- Computer Science
- Crisis Management
- Data Modeling and Analysis
- Design, Analysis, and Manufacturing
- Energy / Transportation
- Imaging
- Innovation
- Leadership / Communication
- Radar
- Real Estate
- Sustainability
- Systems Engineering

NEW Professional Certificate Program in Innovation and Technology

Focused on building competitive advantage through innovation and technology, this credential signifies your depth of knowledge and skill in this high-value area. You can earn your Professional Certificate by completing four on-campus courses.

Learn more about MIT Professional Education and programs.

Visit **professional.mit.edu/techreview** or email **professionaleducation@mit.edu**

Upfront

3 QUESTIONS

Mark Fields

Why has Ford, which you lead as CEO, been more cautious about automated driving technology than many automakers?
We take autonomous driving very, very seriously. And we want to make sure that when we talk about something we have a lot of experience under our belts to inform us. But at the highest level, we are transforming the company from just an auto company to an auto and mobility company and thinking about it in this more holistic way, through what we call Ford Smart Mobility, and one of the elements is autonomous driving.

How do you expect autonomous driving to roll out to the general public, both from Ford and from others?
I think first off we would see fully autonomous vehicles launching in defined areas that have been 3-D mapped. And potentially launched with a service in mind first. Passengers, a passenger kind of ride-sharing service, those are some of the things we're thinking about right now.

Autonomous vehicles so far refuse to break the law—even to avoid an accident. How do we add that aspect of intelligence to driverless cars?
We're going to have to work at it. Let's say you're at a four-way stop. The rule is whoever gets there first and stops can go. But as you know, people don't always follow the law. So one of the things is how would we program into the vehicle that you can actually start creeping forward a bit to signal to the other cars that you're going through? I think this is part of the development process we have in front of us.

—Rachel Metz

A Boost for Solar

Researchers say they've overcome an obstacle to making highly efficient solar devices by combining silicon with a new material.

Silicon probably won't be replaced as the dominant solar material anytime soon, but it might not be too long before it gets a partner from a promising class of materials called perovskites.

A group led by Henry Snaith, a physicist at the University of Oxford and a leading perovskite researcher, has demonstrated what it says is a viable pathway to a device that combines a conventional silicon cell with a perovskite cell to boost the efficiency of the silicon cell by several percentage points.

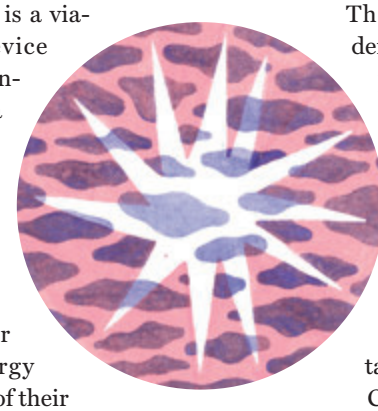

Perovskites, which have captured the interest of solar researchers and energy policy experts because of their rapidly improving performance and low cost, are distinguished by a chemical structure that gives rise to unique electronic properties making them attractive for solar technology. Snaith and his colleagues say the new composition they've developed overcomes a fundamental obstacle to designing a highly efficient device that combines the light-absorbing characteristics of silicon with those of a perovskite material.

The researchers say it should be possible to make a silicon-perovskite "tandem" device that is more than 25 percent efficient—better than today's commercially available silicon cells, which are about 17 to 20 percent efficient.

High-performance tandem devices made of semiconductors other than perovskite have already achieved efficiencies of over 40 percent in the lab, but they are extremely expensive because they require very technically complex manufacturing processes. Making perovskite solar cells is much simpler and cheaper, and the process could be integrated into existing silicon panel manufacturing lines by adding a few steps.

The researchers have only demonstrated the new composition at a small scale, and a lot of work would be needed before it could be available in commercial panels. But a company Snaith cofounded, Oxford PV, is also focused on developing silicon-perovskite tandem devices.

Chris Case, chief technology officer of Oxford PV, says the recent results reflect how quickly researchers are addressing the challenges inherent in making reliable, high-performing tandem cells. Case won't reveal the specifics of his company's technology but says Oxford PV is close to demonstrating full-size devices that are 23 percent efficient and could hit 25 percent before long.

Perovskite-based technologies still face challenges related to the materials' sensitivity to moisture and air, and questions remain about whether perovskite cells can be made durable enough to have the long lifetimes required of power systems. Still, Case says Oxford PV is on track to deliver a commercial product—aimed at silicon panel manufacturers who want to upgrade the efficiency of their products—in 2017. *—Mike Orcutt*

GRAFILU; MONICA RAMOS

#techstyle

the art of innovation in fashion

Museum of Fine Arts Boston

March 6 – July 10

mfa.org/techstyle

Designed by Iris van Herpen and Neri Oxman, printed by Stratasys,
Anthozoa Cape and Skirt, *Voltage* Haute Couture Collection, 2013.
Object Connex multiple-materials; 3-D printed. Museum purchase with funds donated by the Fashion Council. © M. Zoeter x Iris van Herpen. Photography by Ronald Stoops.

Sponsored by the Fashion Council and New Balance

Presented with generous support from The Coby Foundation, Ltd.

Additional support from the David and Roberta Logie Fund for Textile and Fashion Arts and the Consulate General of Israel to New England.

MIT Technology Review

EmTech DIGITAL

Intelligent systems are changing your business. Don't be left behind.

Artificial intelligence already impacts every industry, powering search, social media, and smartphones and tracking personal health and finances. What's coming next promises to be the greatest computing breakthrough of our time. EmTech Digital separates facts from hype – giving you the business intelligence you need.

Learn how to harness new technologies to build or maintain a competitive business advantage.

Gain clear business insights in order to cut through the noise around artificial intelligence and big data so you can chart your strategy.

Meet the sharpest minds leading the next wave of intelligent technologies from institutions including the Allen Institute for AI, Baidu, Carnegie Mellon, Google, IBM, Tesla, and more.

Register Today:

technologyreview.com/emtechdigital

MIT Technology Review subscribers
SAVE 15%.
Enter promotional code
TRSub
at registration.

May 23-24, 2016
St. Regis Hotel
San Francisco, CA

EmTech Digital 2016 examines:

The Technologies Powering Change

The latest research on artificial intelligence techniques, including deep learning and speech and image recognition, that are providing machines with valuable new capabilities and making the automation of more business decisions possible.

The Business Impacts and Opportunities

How AI will continue to change the face of every industry, including:

+ Health care and biomedicine

+ Retail and finance

+ Transportation and auto

+ Media and advertising

... and more

Featured Speakers Include:

Sterling Anderson

Senior Product Manager, Tesla

Delivering on the Promise of Autonomous Vehicles

Maja Matarić

Founding Director, USC Robotics and Autonomous Systems Center

Engineering Social Robots: Next-Generation Human-Robot Interaction

Peter Norvig

Director of Research, Google

State-of-the-Art AI: Building Tomorrow's Intelligent Systems

Andrew Ng

Chief Scientist, Baidu

Deep Learning in Practice: Speech Recognition and Beyond

Manuela Veloso

Herbert A. Simon University Professor, Carnegie Mellon University

Next Generation Human-Machine Collaboration

Q+A

Will Machines Eliminate Us?

Yoshua Bengio leads one of the world's preëminent research groups developing a powerful AI technique known as deep learning. The startling capabilities that deep learning has given computers in recent years, from human-level voice recognition and image classification to basic conversational skills, have prompted warnings about the progress AI is making toward matching, or perhaps surpassing, human intelligence. Prominent figures such as Stephen Hawking and Elon Musk have even cautioned that artificial intelligence could pose an existential threat to humanity. Musk and others are investing millions of dollars in researching the potential dangers of AI, as well as possible solutions. But the direst statements sound overblown to many of the people who are actually developing the technology. Bengio, a professor of computer science at the University of Montreal, put things in perspective in an interview with *MIT Technology Review*'s senior editor for AI and robotics, Will Knight.

Should we worry about how quickly artificial intelligence is advancing?
There are people who are grossly overestimating the progress that has been made. There are many, many years of small progress behind a lot of these things, including mundane things like more data and computer power. The hype isn't about whether the stuff we're doing is useful or not—it is. But people underestimate how much more science needs to be done. And it's difficult to separate the hype from the reality because we are seeing these great things and also, to the naked eye, they look magical.

Is there a risk that AI researchers might accidentally "unleash the demon," as Musk has put it?
It's not like somebody found some magical recipe suddenly. Things are much more complicated than the simple story some people would like to tell. Journalists would sometimes like to tell the story that someone in their garage will have this amazing idea, and then we have a breakthrough and have AI. Similarly, companies want to tell a nice little story that "Oh, we have this revolutionary technology that's going to change the world—AI is almost here, and we are the company that's going to deliver it." That's not at all how it works.

What about the idea, central to these concerns, that AI could somehow start improving itself and then become difficult to control?
It's not how AI is built these days. Machine learning means you have a painstaking, slow process of acquiring information through millions of examples. A machine improves itself, yes, but very, very slowly, and in very specialized ways. And the kinds of algorithms we play with are not at all like little virus things that are self-programming. That's not what we're doing.

KRISTINA COLLANTES

What are some of the big unsolved problems with AI?

Unsupervised learning is really, really important. Right now, the way we're teaching machines to be intelligent is that we have to tell the computer what is an image, even at the pixel level. For autonomous driving, humans label huge numbers of images of cars to show which parts are pedestrians or roads. It's not at all how humans learn, and it's not how animals learn. We're missing something big. This is one of the main things we're doing in my lab, but there are no short-term applications—it's probably not going to be useful to build a product tomorrow.

"We're missing something big. We've been making pretty fast progress, but it's still not at the level where we would say the machine understands. We are still far from that."

Another big challenge is natural language understanding. We've been making pretty fast progress in the past few years, so it's very encouraging. But it's still not at the level where we would say the machine understands. That would be when we could read a paragraph and then ask any question about it, and the machine would basically answer in a reasonable way, as a human would. We are still far from that.

What approaches beyond deep learning will be needed to create a true machine intelligence?

Traditional endeavors, including reasoning and logic—we need to marry these things with deep learning in order to move toward AI. I'm one of the few people who think that machine-learning people, and especially deep-learning people, should pay more attention to neuroscience. Brains work, and we still don't know why in many ways. Improving that understanding has a great potential to help AI research.

And I think that neuroscience people would gain a lot from keeping track of what we do and trying to fit what they observe of the brain with the kinds of concepts we are developing in machine learning.

Did you ever think you'd have to explain to people that AI isn't about to take over the world? That must be odd.

It's certainly a new concern. For so many years, AI has been a disappointment. As researchers we fight to make the machine slightly more intelligent, but they are still so stupid. I used to think we shouldn't call the field artificial intelligence but artificial stupidity. Really, our machines are dumb, and we're just trying to make them less dumb.

Now, because of these advances that people can see with demos, now we can say, "Oh, gosh, it can actually say things in English, it can understand the contents of an image." Well, now we connect these things with all the science fiction we've seen and it's like, "Oh, I'm afraid!"

Okay, but surely it's still important to think now about the eventual consequences of AI.

Absolutely. We ought to be talking about these things. The thing I'm more worried about, in a foreseeable future, is not computers taking over the world. I'm more worried about misuse of AI. Things like bad military uses, manipulating people through really smart advertising; also, the social impact, like many people losing their jobs. Society needs to get together and come up with a collective response, and not leave it to the law of the jungle to sort things out.

Stay ahead of the technologies that matter most to your business

Business Reports

Get a one-year subscription or order individual copies at

technologyreview.com /businessreports

MIT Technology Review

10 Breakthrough Technologies 2016

Which of today's emerging technologies have a chance at solving a big problem and opening up new opportunities? Here are our picks. The 10 on this list all reached an impressive milestone in the past year or are on the verge of reaching one.

JESSICA SVENDSEN

Immune Engineering

Genetically engineered immune cells are saving the lives of cancer patients. That may be just the start.

By Antonio Regalado

The doctors looking at Layla Richards saw a little girl with leukemia bubbling in her veins. She'd had bags and bags of chemotherapy and a bone marrow transplant. But the cancer still thrived. By last June, the 12-month-old was desperately ill. Her parents begged—wasn't there anything?

There was. In a freezer at her hospital—Great Ormond Street, in London—sat a vial of white blood cells. The cells had been genetically altered to hunt and destroy leukemia, but the hospital hadn't yet sought permission to test them. They were the most extensively engineered cells ever proposed as a therapy, with a total of four genetic changes, two of them introduced by the new technique of genome editing.

Soon a doctor from Great Ormond was on the phone to Cellectis, a biotechnology company with French roots that is now located on the East Side of Manhattan. The company owned the cancer treatment, which it had devised using a gene-editing method called TALENs, a way of making cuts and fixes to DNA in living cells. "We got a call. The doctors said, 'We've got a girl who is out of T cells and out of options,'" André Choulika, the CEO of Cellectis, remembers. "They wanted one of the vials made during quality-control testing."

The doctors hoped to make Layla a "special," a patient who got the drug outside a clinical trial. It was a gamble, since the treatment had been tried only in mice. If it failed, the com-

Breakthrough
Killer T cells programmed to wipe out cancer.

Why It Matters
Cancer, multiple sclerosis, and HIV could all be treated by engineering the immune system.

Key Players in Immune Therapies
- Cellectis
- Juno Therapeutics
- Novartis

ILLUSTRATIONS BY JON HAN

pany's stock and reputation could tank, and even if it succeeded, the company might get in trouble with regulators. "It was saving a life versus the chance of bad news," Choulika says.

T cells can crawl, sense things, and even kill other cells. They're little robots.

Cellectis began developing the treatment in 2011 after doctors in New York and Philadelphia reported that they'd found a way to gain control over T cells, the so-called killer cells of the immune system. They had shown that they could take T cells from a person's bloodstream and, using a virus, add new DNA instructions to aim them at the type of blood cell that goes awry in leukemia. The technique has now been tested in more than 300 patients, with spectacular results, often resulting in complete remission. A dozen drug firms and biotechnology companies are now working to bring such a treatment to market.

The T cells created by Cellectis could have even broader applications. The previous treatments use a person's own cells. But some patients, especially small children like Layla, don't have enough T cells.

Foreseeing this problem, Cellectis had set out to use gene editing to create a more highly engineered but ultimately simpler "universal" supply of T cells made from the blood of donors. The company would still add the new DNA, but it would also use gene editing to delete the receptor that T cells normally use to sniff out foreign-looking molecules.

"The T cell has a huge potential for killing. But the thing you can't do is inject T cells from Mr. X into Mr. Y," Choulika says. "They'd recognize Mr. Y as 'non-self' and start firing off at everything, and the patient will melt down." But if the T cells are stripped down with gene editing, like the ones that were sitting in Great Ormond's freezer, that risk is mostly eliminated. Or so everyone hoped.

In November, Great Ormond announced that Layla was cured. The British press jumped on the heartwarming story of a brave kid and daring doctors. Accounts splashed on front pages sent Cellectis's stock price shooting upward. Two weeks later, the drug companies Pfizer and Servier announced they would ante up $40 million to purchase rights to the treatment.

Although many of the details of Layla's case have yet to be disclosed, and some cancer experts say the role of the engineered T cells in her cure remains murky, her recovery pointed a spotlight on "immune engineering," and on the way that advances in controlling and manipulating the immune system are leading to unexpected breakthroughs in cancer treatment. They also could lead to new treatments for HIV and autoimmune diseases like arthritis and multiple sclerosis.

Known killer

The human immune system has been called nature's "weapon of mass destruction." It has a dozen major cell types, including several kinds of T cells. It defends against viruses it's never seen before, suppresses cancer (though not always), and for the most part manages to avoid harming the body's own tissue. It even has a memory, which is the basis of all vaccines.

More than 100 years ago, the American surgeon William Coley observed that an unexpected infection could sometimes make a tumor evaporate. Sub-

Boom Times

Immune-engineering startups have gone public, raising large sums for human trials.

Company	Amount Raised in IPO	Date
Kite Pharma	$134 million	June 2014
Juno Therapeutics	$304 million	December 2014
Bellicum Pharmaceuticals	$160 million	December 2014
Cellectis	$228 million	March 2015

sequently, Coley injected streptococcal cultures into cancer patients and saw the tumors shrink in some cases. The finding, published in 1893, showed the immune system could confront and fight cancer—but how did it work? Until recently, the answers weren't known, and cancer immunotherapy was seen as a failed idea.

But scientists have gradually mapped the network of molecules that govern how the immune system interacts with a tumor. And over the last few years, these insights have allowed drug companies and labs to start tinkering with the immune system's behavior. "From 40 years and more of science, we know the general nature of the conversation between the tumor cells and the immune system," says Philip Sharp, a biologist at MIT's Koch Institute for Integrative Cancer Research and a recipient of the 1993 Nobel Prize in medicine. "That's the conversation we're trying to join in order to have a therapeutic effect. We are still at the level of a five-year-old kid. We know there are nouns, and that there are verbs. But the diversity of the vocabulary is still being mapped out."

The most extreme of these proposals is to change the genetic instructions inside the T cell itself, something that's become much easier using gene-editing methods like TALENs and the even newer CRISPR. Last year, the gene-editing startups Editas Medicine and Intellia Therapeutics each struck deals with companies developing T-cell-based therapeutics. "It's the perfect setup," says Jeffrey Bluestone, a researcher at the University of California, San Francisco. "Immune cells are machines that work pretty well, but we can make them work even better."

A Time Line of Engineering Immunity

500 million years ago
Jawed fish are first to develop "adaptive" immunity—specialized cells that learn from, remember, and respond to threats.

1796
Edward Jenner inoculates a boy against smallpox using pus from a cowpox blister. It is celebrated as the first vaccine.

1893
New York surgeon William Coley believes cancer can be cured by an immune response. He uses live bacteria, called Coley's toxins, to treat tumors.

1908
German doctor Paul Ehrlich wins a Nobel Prize for his theories about the immune system. He introduces the idea of the "Wundermittel," or magic bullet—the precursor of today's targeted drugs.

1971
President Nixon declares a "War on Cancer." The National Cancer Institute's budget rises to $378 million, or $2.1 billion in current dollars. Today it is $4.95 billion.

1981
The HIV epidemic begins. By 1987 the first antiretroviral drug treatment, AZT, goes on sale. A vaccine remains elusive to this day.

1983–1987
Scientists discover the T cell antigen receptor. It is what killer T cells use to identify virus-infected cells and cancer.

2000
Two immune-deficient children are cured in France of "bubble boy" disease in the first successful use of gene therapy. A missing gene is added to their bone marrow.

2011
The first immune checkpoint inhibitor, ipilimumab, is approved in the United States to treat late-stage melanoma. The drug unleashes T cells. In many patients, the results are dramatic.

2011
Carl June of the University of Pennsylvania reports the successful treatment of leukemia using genetically modified T cells.

2015
Former U.S. president Jimmy Carter, at 90, receives immune therapy treatments for melanoma and a brain cancer. His brain scans are later clear.

2016
Recognizing "amazing advances" in immune therapy, President Obama and Vice President Biden announce a new "moon shot" to cure cancer.

Researchers are building on decades of research (and several Nobel Prizes involving immunology) that worked out many important details, including how T cells recognize invaders and go in for the kill. Seen through a microscope, these cells display almost animal-like behavior: they crawl, probe, then grab another cell and shoot it full of toxic granules. "What's exciting is they have the ability to move all around; they're autonomous," says Wendell Lim, a synthetic biologist who is also at UCSF. "Immune cells talk to other cells, they deliver poisons, they can change what happens in a microenvironment, they have a memory, and they make more of themselves. I think of them as little robots."

Lim is now breaking new ground in what he calls "synthetic immunology." This year and last, he produced some futuristic T cells. Tested only in mice so far, the cells deploy their targeted search-and-kill behavior only if a specific drug is added—a feature that could be used to turn the cells on at specific places and times, which Lim calls "remote control." Another T cell he designed is a two-stage affair, which kills only if it locates not one but two different markers on a cancer cell; it is like a dual authentication method for the enemy cell. Lim thinks of it as a sensing circuit or "advanced Google search."

Such work is critical because targeting T cells to tumors of the liver, lung, or brain is dangerous, and some patients have been killed in trials. The problem has been friendly fire. So far, easy ways to target only cancer cells are lacking. Lim has founded his own startup, Cell Design Labs, to commercialize his engineering ideas. He declined to say how much money he has raised, but he says everyone working with T cells is stunned by the kind of money being thrown at the idea. "It's a 'wow' type of situation," he says.

Googling cures

The search to expand immune therapy now involves not only the world's largest drug companies but also tech firms. Sharp says that last year Google held two summits at MIT of top immune oncologists and bioengineers to determine what parts of the problem could be "Googlified." Attendees say the search giant paid special attention to new research techniques that fingerprint cells from a tumor biopsy in rapid-fire fashion. These methods might generate big data about what immune system cells are actually doing inside a tumor, and new clues about how to influence them. So far, Google's life science unit, named Verily, hasn't revealed its plans in cancer immunotherapy. But in New York's Union Square, I met Jeffrey Hammerbacher, a former Facebook employee who now runs a lab that is part of Mount Sinai, the hospital and medical school. With 12 programmers in a light-soaked loft—the nearest thing to blood and guts is a photo of an exhausted surgeon on the wall—he's also spending time on T cells. He's developing software to

Big Deals

T-cell companies have sought agreements with drug companies and specialists in gene editing.

August 2012
Swiss drug giant Novartis forms a sweeping alliance with the University of Pennsylvania, site of early successes using engineered T cells.

January 2015
Novartis buys CRISPR gene-editing rights from Intellia Therapeutics. Juno and Editas Medicine later strike a similar deal for $25 million.

June 2015
Biotech firm Celgene pays Seattle-based Juno $1 billion for a slice of its T-cell treatment portfolio.

November 2015
Drug companies Pfizer and Servier pay Cellectis $40 million for rights to the first "off the shelf" T-cell treatment for leukemia.

January 2016
Food maker Nestlé pays $120 million to a startup named Seres for bacteria pills able to ward off infection and immune disorders.

January 2016
Juno pays $125 million to buy AbVitro, a Boston company that can sequence the DNA inside individual T cells.

interpret the DNA sequence in a patient's cancer and predict from it how to goose the response of killer T cells.

A clinical trial by Mount Sinai should start this year. The patients receive a dose of abnormal protein fragments that Hammerbacher's software predicts will train T cells to attack the cancer. "What was fun was that what we submitted to the [U.S. Food and Drug Administration] was not a molecule but an algorithm," he says. "It might be one of the first times the output of a program is the therapy."

> "Where the technology stands, it's a pretty radical treatment."

In January, Juno Therapeutics (see "Biotech's Coming Cancer Cure," July/August 2015) paid $125 million to acquire AbVitro, a Boston-area company that specializes in sequencing the DNA inside single T cells. Now Juno is trying to locate T cells that are active inside cancers and study their receptors. Juno's chief scientist, Hyam Levitsky, says an experiment that used to take seven months now takes seven days. And data is piling up: an average experiment generates 100 gigabytes of information. "A lot of what is happening is technology-driven," he says. "The questions have been there for a while, but there was no way to get at the answers. Now we're visualizing them with new technology in ways we never could before."

Beyond cancer

In March Pfizer appointed John Lin to head its San Francisco biotech unit, which develops cancer drugs and recently started making engineered T cells. He says the company had been negotiating with Cellectis well before the news of Layla's treatment and that no one there was even aware the girl had been treated before it hit the news. "The publicity was a big surprise," he says.

Lin says years of scientific work have finally resulted in a level of mastery that makes therapeutic products seem practical. He thinks the treatments will go beyond leukemia, and beyond cancer. "We think that this fundamental principle, engineering human cells, could have broad implications," he says, "and the immune system will be the most convenient vehicle for it, because they can move and migrate and play such important roles."

Researchers are already working on autoimmune disorders, like diabetes, multiple sclerosis, and lupus. Infectious disease is also in the sights of T-cell engineers. Edward Berger, a virologist at the National Institutes of Health who helped discover how HIV enters human cells, thinks it may be possible to permanently keep the virus in check, a so-called "functional cure." In February, he says, he will start giving monkeys T cells genetically programmed to find and destroy any cell in which the simian version of HIV is replicating.

The actual process isn't as simple as the theory. Berger is sure that years of missteps and do-overs lie ahead. Also, most protocols involving engineered T cells require patients, or monkeys, to take drugs that temporarily kill off their own T cells, which isn't without risks. "Where the technology stands, it's a pretty radical treatment," Berger says. "You aren't going to use it on a cold sore." But despite all the progress that has been made treating HIV, a better approach is still needed. Because the virus hides in the body even after treatment, patients have to take antiretroviral drugs for life. With immune engineering, maybe not. Berger sees the chance of a one-time treatment that can hold the virus in check for good.

"I was totally inspired by the cancer work," he says. "They cured leukemia, and we've borrowed it from them. The extension of those ideas for engineering the immune system against other things that ail people is a major front. I think HIV is the best candidate in infectious disease. If you talk to the HIV community, they are crying for a cure—a treatment that, ideally, you do once and never again."

Precise Gene Editing in Plants

CRISPR offers an easy, exact way to alter genes to create traits such as disease resistance and drought tolerance.

By David Talbot

Breakthrough
The ability to cheaply and precisely edit plant genomes without leaving foreign DNA behind.

Why It Matters
We need to increase agricultural productivity to feed the world's growing population, which is expected to reach 10 billion by 2050.

Key Players in Engineering Crops
- The Sainsbury Laboratory and John Innes Centre, Norwich, U.K.
- Seoul National University
- University of Minnesota
- Institute of Genetics and Developmental Biology, Beijing

A new gene-editing method is providing a precise way to modify crops in hopes of making them yield more food and resist drought and disease more effectively. Research in the past year has shown that the resulting plants have no traces of foreign DNA, making it possible that they will not fall under existing regulations governing genetically modified organisms and will sidestep many of the consumer concerns over these GMOs.

The technology is known as CRISPR (see "10 Breakthrough Technologies 2014: Genome Editing"), and plants modified with it are sprouting in laboratory greenhouses around the world. Already, a lab in China has used it to create a fungus-resistant wheat; several groups in China are using the technique on rice in efforts to boost yields; and a group in the U.K. has used it to tweak a gene in barley that helps govern seed germination, which could aid efforts to produce drought-resistant varieties. Indeed, because it's so easy to do and the plants could avoid the lengthy and expensive regulatory process associated with GMOs, the method is increasingly being used by research labs, small companies, and public plant breeders unwilling to take on the expense and risks of conventional genetic engineering.

The gene-editing technique could be critical in helping scientists keep up with the constantly evolving microbes that attack crops, says Sophien Kamoun, who leads a research group at the Sainsbury Lab in Norwich, England, that is applying the technology to potatoes, tomatoes, and other crops to fight fungal diseases. "It takes millions of dollars and many years of work to go through the regulatory process," Kamoun says. "But the pathogens don't sit and wait for you; they keep evolving and changing."

A version of CRISPR he co-developed paved the way for recent work on barley and a broccoli-like plant at the John Innes Centre, a plant science research center also in Norwich. Kamoun and colleagues showed that the second generation of some of the edited plants contain none of the foreign DNA that had been used to create the first generation. (Though CRISPR doesn't require inserting foreign genes, it does typically use bits of bacterial genetic material to target the editing.) Meanwhile, a group at Seoul National University has avoided leaving any foreign genetic material even in first-generation plants.

Big and small companies alike are jumping in. DuPont Pioneer has already invested in Caribou Biosciences, the CRISPR startup cofounded by Jennifer Doudna, one of the inventors of the technology, and is using it in experiments on corn, soybeans, wheat, and rice. It hopes to sell seeds bred with CRISPR technology in as little as five years.

The big question is whether CRISPR crops will be governed by the same regulations as GMOs. The U.S. Department of Agriculture has already said some examples of gene-edited corn, potatoes, and soybeans (edited using a different method, known as TALENs) don't fall under existing regulations. But both the United States and the more restrictive European Union are now conducting reviews of today's regulations. And Chinese authorities have not said whether they will allow the crops to be planted.

GRANT CORNETT

Conversational Interfaces

Powerful speech technology from China's leading Internet company makes it much easier to use a smartphone.

By Will Knight

Breakthrough
Combining voice recognition and natural language understanding to create effective speech interfaces for the world's largest Internet market.

Why It Matters
It can be time-consuming and frustrating to interact with computers by typing.

Key Players in Voice Recognition and Language Processing
- Baidu
- Google
- Apple
- Nuance
- Facebook

Stroll through Sanlitun, a bustling neighborhood in Beijing filled with tourists, karaoke bars, and luxury shops, and you'll see plenty of people using the latest smartphones from Apple, Samsung, or Xiaomi. Look closely, however, and you might notice some of them ignoring the touch screens on these devices in favor of something much more efficient and intuitive: their voice.

A growing number of China's 691 million smartphone users now regularly dispense with swipes, taps, and tiny keyboards when looking things up on the country's most popular search engine, Baidu. China is an ideal place for voice interfaces to take off, because Chinese characters were hardly designed with tiny touch screens in mind. But people everywhere should benefit as Baidu advances speech technology and

TOMI UM

makes voice interfaces more practical and useful. That could make it easier for anyone to communicate with the machines around us.

"I see speech approaching a point where it could become so reliable that you can just use it and not even think about it," says Andrew Ng, Baidu's chief scientist and an associate professor at Stanford University. "The best technology is often invisible, and as speech recognition becomes more reliable, I hope it will disappear into the background."

Voice interfaces have been a dream of technologists (not to mention science fiction writers) for many decades. But in recent years, thanks to some impressive advances in machine learning, voice control has become a lot more practical.

The systems offer a glimpse of a future in which there's less need to learn a new interface for every device.

No longer limited to just a small set of predetermined commands, it now works even in a noisy environment like the streets of Beijing or when you're speaking across a room. Voice-operated virtual assistants such as Apple's Siri, Microsoft's Cortana, and Google Now come bundled with most smartphones, and newer devices, like Amazon's Alexa, offer a simple way to look up information, cue up songs, and build shopping lists with your voice. These systems are hardly perfect, sometimes mishearing and misinterpreting commands in comedic fashion, but they are improving steadily, and they offer a glimpse of a graceful future in which there's less need to learn a new interface for every new device.

Baidu is making particularly impressive progress, especially with the accuracy of its voice recognition, and it has the scale to advance conversational interfaces even further. The company—founded in 2000 as China's answer to Google, which is currently blocked there—dominates the country's domestic search market, with 70 percent of all queries. And it has evolved into a purveyor of many services, from music and movie streaming to banking and insurance.

A more efficient mobile interface would come as a big help in China. Smartphones are far more common than desktops or laptops, and yet browsing the Web, sending messages, and doing other tasks can be painfully slow and frustrating. There are thousands of Chinese characters, and although a system called Pinyin allows them to be generated phonetically from Latin ones, many people (especially those over 50) do not know the system. It's also common in China to use messaging apps such as WeChat to do all sorts of tasks, such as paying restaurant tabs. And yet in many of China's poorer regions, where there is perhaps more opportunity for the Internet to have big social and economic effects, literacy levels are still low.

"It is a challenge and an opportunity," says Ng, who was named one of *MIT Technology Review*'s Innovators Under 35 in 2008 for his work in AI and robotics at Stanford. "Rather than having to train people used to desktop computers to new behaviors appropriate for cell phones, many of them can learn the best ways to use a mobile device from the start."

Ng believes that voice may soon be reliable enough to be used for interacting with all sorts of devices. Robots or home appliances, for example, could be easier to deal with if you could simply talk to them. The company has research teams at its headquarters in Beijing and at a facility in Silicon Valley

Booming smartphone market
Innovation in voice technologies has been driven by the surge of mobile Internet use in China.

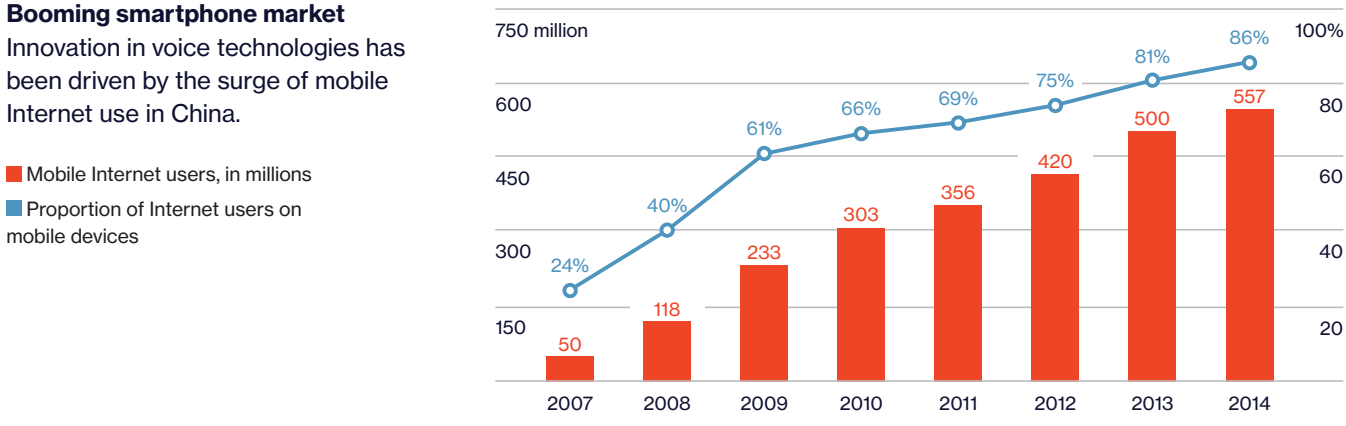

SOURCE: CHINA INTERNET NETWORK INFORMATION CENTER

Researchers at Baidu's headquarters in Beijing are plugging away at a digital assistant that can hold a conversation.

QILAI SHEN/BLOOMBERG VIA GETTY IMAGES

Few of those behind Deep Speech 2 speak Mandarin or Cantonese. It's a universal language engine.

that are dedicated to advancing the accuracy of speech recognition and working to make computers better at parsing the meaning of sentences.

Jim Glass, a senior research scientist at MIT who has been working on voice technology for the past few decades, agrees that the timing may finally be right for voice control. "Speech has reached a tipping point in our society," he says. "In my experience, when people can talk to a device rather than via a remote control, they want to do that."

Last November, Baidu reached an important landmark with its voice technology, announcing that its Silicon Valley lab had developed a powerful new speech recognition engine called Deep Speech 2. It consists of a very large, or "deep," neural network that learns to associate sounds with words and phrases as it is fed millions of examples of transcribed speech. Deep Speech 2 can recognize spoken words with stunning accuracy. In fact, the researchers found that it can sometimes transcribe snippets of Mandarin speech more accurately than a person.

Baidu's progress is all the more impressive because Mandarin is phonetically complex and uses tones that transform the meaning of a word. Deep Speech 2 is also striking because few of the researchers in the California lab where the technology was developed speak Mandarin, Cantonese, or any other variant of Chinese. The engine essentially works as a universal speech system, learning English just as well when fed enough examples.

Most of the voice commands that Baidu's search engine hears today are simple queries—concerning tomorrow's weather or pollution levels, for example. For these, the system is usually impressively accurate. Increasingly, however, users are asking more complicated questions. To take them on, last year the company launched its own voice assistant, called DuEr, as part of its main mobile app. DuEr can help users find movie show times or book a table at a restaurant.

The big challenge for Baidu will be teaching its AI systems to understand and respond intelligently to more complicated spoken phrases. Eventually, Baidu would like for DuEr to take part in a meaningful back-and-forth conversation, incorporating changing information into the discussion. To get there, a research group at Baidu's Beijing offices is devoted to improving the system that interprets users' queries. This involves using the kind of neural-network technology that Baidu has applied in voice recognition, but it also requires other tricks. And Baidu has hired a team to analyze the queries fed to DuEr and correct mistakes, thus gradually training the system to perform better.

"In the future, I would love for us to be able to talk to all of our devices and have them understand us," Ng says. "I hope to someday have grandchildren who are mystified at how, back in 2016, if you were to say 'Hi' to your microwave oven, it would rudely sit there and ignore you."

Reusable Rockets

Rockets typically are destroyed on their maiden voyage. But now they can make an upright landing and be refueled for another trip, setting the stage for a new era in spaceflight.

By Brian Bergstein

Breakthrough
Rockets that can launch payloads into orbit and then land safely.

Why It Matters
Lowering the cost of flight would open the door to many new endeavors in space.

Key Players in the New Space Industry
- SpaceX
- Blue Origin
- United Launch Alliance

Thousands of rockets have flown into space, but not until 2015 did one return like this: it came down upright on a landing pad, steadily firing to control its descent, almost as if a movie of its launch were being played backward. If this can be done regularly and rockets can be refueled over and over, spaceflight could become a hundred times cheaper.

Two tech billionaires made it happen. Jeff Bezos's Blue Origin first pulled off a landing in November; Elon Musk's SpaceX did it in December. The companies are quite different—Blue Origin hopes to propel tourists in capsules on four-minute space rides, while SpaceX already launches satellites and space station supply missions—but both need reusable rockets to improve the economics of spaceflight.

Blasting things into space has been expensive because rockets cost tens of millions of dollars and fly once before burning up in a free fall back through the atmosphere. SpaceX and Blue Origin instead bring theirs down on fold-out legs, a trick that requires onboard software to fire thrusters and manipulate flaps that slow or nudge the rockets at precise moments.

SpaceX has the harder job because Blue Origin's craft go half as fast and half as high and stay mostly vertical, whereas SpaceX's rockets have to switch out of a horizontal position. A reminder of how many things can go wrong came in January, when SpaceX just missed a second landing because a rocket leg didn't latch into place. Even so, it's now clear that the future of spaceflight will be far more interesting than the Apollo-era hangover of the past 40 years.

THIS PAGE COURTESY OF EMTECH; OPPOSITE PAGE COURTESY OF SPACEX

Facing page: SpaceX made test landings in Texas. This page: A long exposure captured a SpaceX rocket taking off and returning to Cape Canaveral, Florida.

Robots That Teach Each Other

What if robots could figure out more things on their own and share that knowledge among themselves?

By Amanda Schaffer

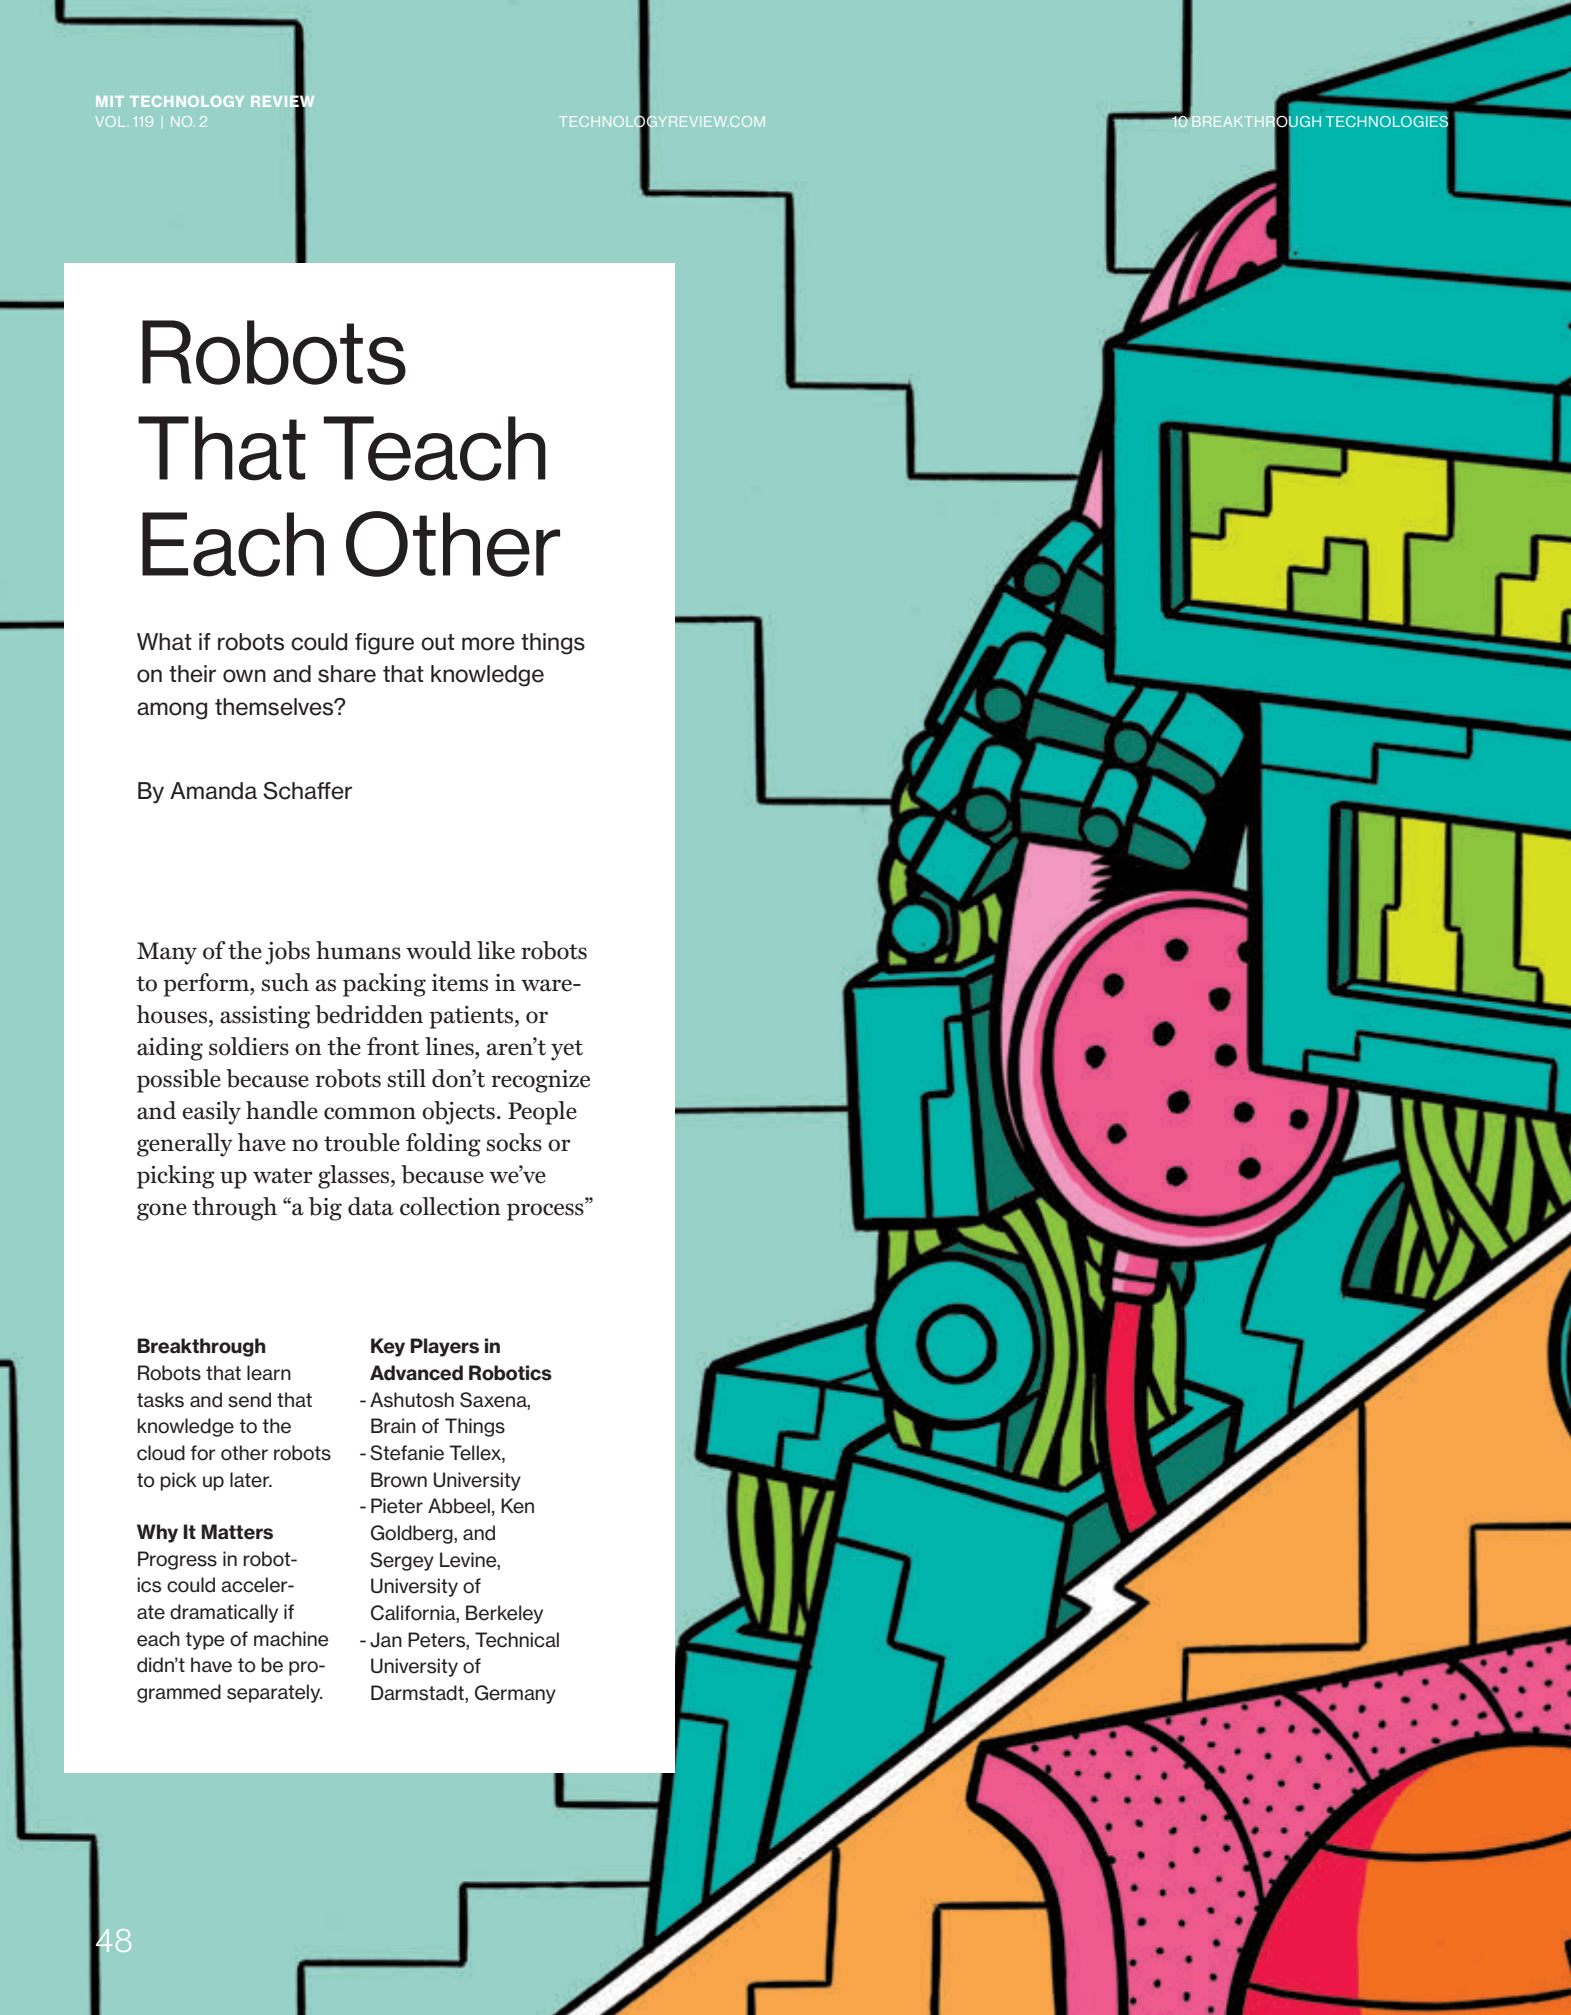

Many of the jobs humans would like robots to perform, such as packing items in warehouses, assisting bedridden patients, or aiding soldiers on the front lines, aren't yet possible because robots still don't recognize and easily handle common objects. People generally have no trouble folding socks or picking up water glasses, because we've gone through "a big data collection process"

Breakthrough
Robots that learn tasks and send that knowledge to the cloud for other robots to pick up later.

Why It Matters
Progress in robotics could accelerate dramatically if each type of machine didn't have to be programmed separately.

Key Players in Advanced Robotics
- Ashutosh Saxena, Brain of Things
- Stefanie Tellex, Brown University
- Pieter Abbeel, Ken Goldberg, and Sergey Levine, University of California, Berkeley
- Jan Peters, Technical University of Darmstadt, Germany

KRISTIAN HAMMERSTAD

One researcher expects "an explosion in the ability of robots."

Stefanie Tellex and a Baxter robot at Brown University.

called childhood, says Stefanie Tellex, a computer science professor at Brown University. For robots to do the same types of routine tasks, they also need access to reams of data on how to grasp and manipulate objects. Where does that data come from? Typically it has come from painstaking programming. But ideally, robots could get some information from each other.

That's the theory behind Tellex's "Million Object Challenge." The goal is for research robots around the world to learn how to spot and handle simple items from bowls to bananas, upload their data to the cloud, and allow other robots to analyze and use the information.

Tellex's lab in Providence, Rhode Island, has the air of a playful preschool. On the day I visit, a Baxter robot, an industrial machine produced by Rethink Robotics, stands among oversized blocks, scanning a small hairbrush. It moves its right arm noisily back and forth above the object, taking multiple pictures with its camera and measuring depth with an infrared sensor. Then, with its two-pronged gripper, it tries different grasps that might allow it to lift the brush. Once it has the object in the air, it shakes it to make sure the grip is secure. If so, the robot has learned how to pick up one more thing.

The robot can work around the clock, frequently with a different object in each of its grippers. Tellex and her graduate student John Oberlin have gathered—and are now sharing—data on roughly 200 items, starting with such things as a child's shoe, a plastic boat, a rubber duck, a garlic press and other cookware, and a sippy cup that originally belonged to her three-year-old son. Other scientists can contribute their robots' own data, and Tellex hopes that together they will build up a library of information on how robots should handle a million different items.

PHOTOGRAPHS BY LEONARD GRECO

Each time the robot determines the best way to grasp and hold something, it files that data away in a format other robots can use.

Eventually, robots confronting a crowded shelf will be able to "identify the pen in front of them *and* pick it up," Tellex says.

Projects like this are possible because many research robots use the same standard framework for programming, known as ROS. Once one machine learns a given task, it can pass the data on to others—and those machines can upload feedback that will in turn refine the instructions given to subsequent machines. Tellex says the data about how to recognize and grasp any given object can be compressed to just five to 10 megabytes, about the size of a song in your music library.

Tellex was an early partner in a project called RoboBrain, which demonstrated how one robot could learn from another's experience. Her collaborator Ashutosh Saxena, then at Cornell, taught his PR2 robot to lift small cups and position them on a table. Then, at Brown, Tellex downloaded that information from the cloud and used it to train her Baxter, which is physically different, to perform the same task in a different environment.

Such progress might seem incremental now, but in the next five to 10 years, we can expect to see "an explosion in the ability of robots," says Saxena, now CEO of a startup called Brain of Things. As more researchers contribute to and refine cloud-based knowledge, he says, "robots should have access to all the information they need, at their fingertips."


Photos
Camera
Weather
Clock
Maps
Videos
Notes
Reminders
Stocks
Game Center
Newsstand
iTunes Store
App Store
iBooks
Health
Passbook
Settings

DNA App Store

An online store for information about your genes will make it cheap and easy to learn more about your health risks and predispositions.

By Antonio Regalado

While driving and listening to National Public Radio one day, Justin Kao heard about the discovery of a "sweet tooth gene" that makes you more likely to crave sweets. "Oh my God," thought Kao, who has always loved cookies. "I would pay $5 to know if I had that."

Kao is hoping that millions of other people will be just as eager to spend a few bucks for tidbits revealed in their DNA. He is a cofounder of Helix, a San Francisco–based company that last summer secured more than $100 million in a quest to create the first "app store" for genetic information.

Our genomes hold information about our health risks, our physical traits, and whom we're related to. Yet aside from ancestry tests that provide a limited genetic snapshot, there's not a mass market for DNA data. Helix is a bet by Kao's former employer, the buyout firm Warburg Pincus, and Illumina, the leading manufacturer of ultrafast DNA sequencing machines, that what's been missing is the right business model.

Helix's idea is to collect a spit sample from anyone who buys a DNA app, sequence and analyze the customers' genes, and then digitize the findings so they can be accessed by software developers who want to sell other apps. Helix calls the idea "sequence once, query often." (The company says customers will find these apps on websites and possibly in the Android and Apple app stores.)

With its ties to Illumina, Helix thinks it can decode the most important part of a person's genome—all 20,000 genes and a few other bits—at a cost of about $100, about one-fifth of what it costs other companies. That's why Helix can afford its second gambit: to generate and store this type of data for all customers, even if they initially make only one specific genetic query—such as whether they have the sweet tooth gene or a risk for a certain disease. Maybe two guys in a garage will write a $10 app that shows you how old you'll look in 10 years, or which celebrity you are most closely related to. Kao says the tactic will make genetic information available to consumers "at an unprecedentedly low entry price."

The engine to power the app store is being assembled a mile from Illumina's San Diego headquarters, in a building where workmen were still bending sheet metal and laying floor tiles in January. Several miles of data cables strung through the ceiling will be connected to a large farm of sequencing machines, able to process the DNA from a million samples a year. Illumina's CEO, Jay Flatley, also chairman of Helix, has said it could be the largest sequencing center anywhere.

Helix plans to launch the store this year or next. Customers will control their data by deciding who sees it. There's even a "nuclear button" to erase every A, G, C, and T. But key details are still being sorted out. Will people be able to download their DNA information and take it elsewhere? Probably, though they might pay extra for the privilege.

One company working with Helix is Good Start Genetics, a startup in Cambridge, Massachusetts, that offers preconception testing. These DNA tests tell parents-to-be if they share a risk for passing on a serious genetic condition, such as cystic fibrosis. Jeffrey Luber, Good Start's head of business development, says it hopes to reach a larger audience with an app that can report a few important risks. As with browsing on Amazon, he thinks, people will discover things they "didn't know they needed but that [are] targeted to them, and that they want."

A looming question mark is the U.S. Food and Drug Administration, which has kept close tabs on gene tests and will decide how much information Helix apps can reveal. Right now, says Keith Stewart, director of the Center for Individualized Medicine at the Mayo Clinic, most apps that return real medical information—your chance of cancer, say, not just how much Neanderthal is in your DNA—would need agency approval, or at least a doctor in the loop.

"The bottom line is going to be: What are the regulatory constraints on information that is truly useful?" says Mirza Cifric, CEO of Veritas Genomics. His company has been offering since last fall to sequence a person's entire genome and is creating its own app to explore the data, complete with a button to get a FaceTime appointment with a genetic counselor. Cifric hasn't decided whether to create an app with Helix, but he says he shares its core belief: "The genome is an asset that you have for life, and you'll keep going back to it."

Breakthrough
A new business model for DNA sequencing that will make genetic information widely accessible online.

Why It Matters
Your genome determines a great deal about you, including your likelihood of getting certain diseases.

Key Players in Consumer Genomics
- Helix
- Illumina
- Veritas Genomics

JAVIER JAEN

The massive solar manufacturing facility, shown here in late December, is scheduled to begin full-scale production sometime next year.

SolarCity's Gigafactory

A $750 million solar facility in Buffalo will produce a gigawatt of high-efficiency solar panels per year and make residential panels far more attractive to homeowners.

By Richard Martin
Photographs by Gus Powell

In an industrial park near the shore of Lake Erie, hard by the Buffalo River, the future of the solar power industry is under construction. SolarCity's sprawling Buffalo factory, built and paid for by the state of New York, is nearing completion and will soon begin producing some of the most efficient solar panels available commercially. Capable of making 10,000 solar panels a day, or one gigawatt of solar capacity a year, it will be the largest solar manufacturing plant in North America and one of the biggest in the world.

When production begins, SolarCity, already the leading installer of residential solar panels in the United States, will become a vertically integrated manufacturer and provider—doing everything from making the solar cells to putting them on rooftops. At a time when conventional silicon-based solar panels from China have never been cheaper, investing in a new

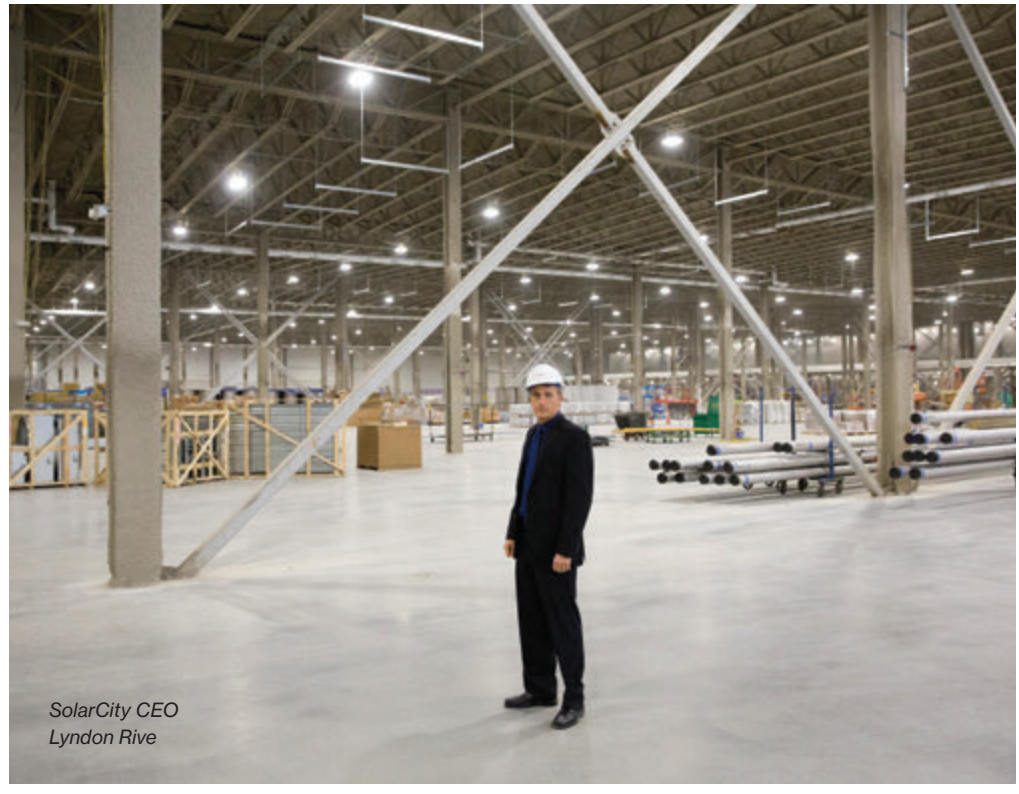

SolarCity CEO Lyndon Rive

SolarCity's new panels use a novel combination of materials.

Breakthrough
Highly efficient solar panels made using a simplified, low-cost manufacturing process.

Why It Matters
The solar industry needs cheaper and more efficient technology to be more competitive with fossil fuels.

Key Players in Photovoltaics
- SolarCity
- SunPower
- Panasonic

The gigafactory covers some 27 acres, making it the largest solar production facility in North America.

type of solar technology is a risky undertaking. However, the potential benefits are huge. The new factory, says SolarCity chief technology officer Peter Rive, could transform both SolarCity's business, which has consistently lost money, and the economics of residential solar power.

Solar panels installed by SolarCity cost the company $2.84 per watt (including sales and marketing plus overhead, in addition to the cost of the hardware), down from $4.73 in 2012. The combination of the new, highly efficient panels, the volume of product coming out of the new factory, and a simplified manufacturing process is a big reason why the company expects its costs for residential solar to fall well below $2.50 per watt by the end of 2017, when the Buffalo facility reaches full production.

Bolstered by federal solar subsidies and "net metering," the rules that allow homeowners to sell excess power back to the grid at retail prices in many states, SolarCity is already leading the way in making residential systems financially attractive to many households, spurring an explosion in the popularity of the rooftop panels. The drop in installed costs could make residential solar even more popular.

"Right now we can sell you energy in 14 states at a rate lower than what you're currently paying the utility," says Rive. The Buffalo factory, he adds, "sets us up for a future where solar plus batteries is cheaper than fossil fuels."

Key to the company's ambitions is a technology it acquired when it bought a small solar company called Silevo in 2014. That technology, which allows it to make panels that are highly efficient at converting sunlight into electricity, traces its origins to the Australian solar power pioneer Martin Green in the late 1970s. It combines a standard crystalline-silicon solar cell with elements of a thin-film cell, along with a layer of a semiconductor oxide. Last October, SolarCity announced that test panels made at a small facility in Fre-

The factory is situated on a former Republic Steel manufacturing site, not far from downtown Buffalo.

mont, California, had tested at just over 22 percent efficiency. Today's commodity silicon-based solar panels have efficiencies of between 16 and 18 percent. SolarCity competitor SunPower previously led the market with cells that can reach 21.5 percent.

Efficiency matters because the panels themselves represent only 15 to 20 percent of the cost of the full installation. Much of the rest comes in what's known as balance-of-system costs: inverters to connect to the grid, materials to house the array, nuts and bolts to attach it to the roof, the labor to install it, and so on. SolarCity's installation, says the company, will require one-third fewer panels to produce the same amount of electricity as conventional installations. "Fewer panels means fewer bits and pieces, less wire, less days on the roof to install," says Francis O'Sullivan, the director of research and analysis at the MIT Energy Initiative.

SolarCity uses a deposition manufacturing process that reduces the number of steps required to make the cells from two dozen or more to just six. It also replaces silver, one of the most expensive elements of conventional solar cells, with less expensive copper.

But the difference in performance between solar panels produced in a small facility like SolarCity's Fremont plant and in a large factory like the Buffalo one could be significant. And scaling up production could be particularly tricky given SolarCity's lack of manufacturing experience. Rive acknowledges that there could be "small risks around the actual time line" in getting the products coming out of Buffalo to match the efficiencies achieved at small scale. Already, SolarCity has pushed back the target date for full production from the Buffalo factory from the first quarter of 2017 to later in the year.

But the real risk lies in the rapid advance of solar technology: a record-setting panel today might look relatively inefficient three or five years down the road. Soon after SolarCity showed off its high-efficiency panels last October, Panasonic topped its rival by claiming that its new panels would reach efficiencies of 22.5 percent. Meanwhile, efficiencies in the lab are even higher: researchers have made exotic solar-cell materials with efficiencies of up to 40 percent. "I think that within 10 years, most manufacturers will be producing panels over 20 percent efficiency, with the best commercial panels reaching over 23 percent," Green says.

O'Sullivan adds: "For now, SolarCity is moving the boat out as far as it can with, generically speaking, contemporary technology. But we're beginning to approach a choke point for the economics of any silicon-based technology"—including the new cells SolarCity is bringing online. Future advances, he says, will entail much lighter, flexible panels that offer much higher efficiencies and are even cheaper to install—and thus produce electricity at a much lower cost.

At that point, the solar panels coming out of the gigafactory may seem as conventional as commodity panels produced in China today. It is, however, SolarCity's willingness to take on such risks that makes the Buffalo facility so ambitious. Over the last 10 years, the Silicon Valley company has made residential solar a popular choice for many consumers through smart marketing and attractive financing. Now it wants to transform solar manufacturing. Whether SolarCity succeeds or fails, it is once again pushing the possibilities of solar power.

COURTESY OF SOLARCITY

Senator Charles Schumer speaks alongside CEO Rive at the Buffalo facility, a key part of New York's effort to revitalize manufacturing in the city.

Proportion of online time spent on mobile devices in the U.S.

TODAY
51%

2008
12%

Proportion of U.S. workers who telecommute

TODAY
37%

1995
9%

Estimated average number of e-mails sent and received by business users each day

TODAY
122

2011
105

Slack

A service built for the era of mobile phones and short text messages is changing the workplace.

By Lee Gomes

Breakthrough
Easy-to-use communication software that is supplanting e-mail as a method of getting work done.

Why It Matters
In many kinds of workplaces, the "water cooler" effect that lets people overhear their colleagues' conversations can enhance productivity.

Key Players in Communication Software
- Slack
- Quip
- Hipchat
- Microsoft

The intra-office messaging system known as Slack is often described as the fastest-growing workplace software the world has ever seen. It surpassed two million daily users less than three years after its launch in 2013.

But what, exactly, makes it so popular?

Slack gives you a centralized place to communicate with your colleagues through instant messages and in chat rooms, which can reduce the time you have to spend on e-mail. Whether you're on a mobile device or a desktop computer, you can upload files, get and manipulate information stored in spreadsheets or other business applications, and easily search through past conversations. But many of the core features have been around since the 1990s. And there have been other "Facebook for the office" software packages that resemble Slack and have failed to generate anything close to the same level of enthusiasm.

The reason for its success lies in part with big trends: more and more people now get work done on mobile devices, in collaboration with people who aren't always in the same office at the same time. But Slack's specific design choices have also been important. Gerald C. Kane, associate professor of information systems at Boston College's Carroll School of Management, points out that Slack funnels messages into streams that everyone who works together can see. That "allows you to 'overhear' what is going on in an organization, which research has shown can lead to business impact," he says. "It's a kind of ambient awareness that you just don't get from e-mail."

Kristina Lerman, a specialist in social computing at the Information Sciences Institute at the University of Southern California, notes that Slack messages tend to be short and casual, much more like the mobile text messages that people are increasingly favoring over e-mail in their personal life. This creates the perception that keeping in touch with coworkers is effortless. "You get the feeling that you are quickly responding to everything that is happening around you," Lerman says.

In fact, Slack makes it so easy to create messages that it might end up placing as many demands on people's time as e-mail traditionally has, albeit with a hip and friendly interface. "There are limits to the amount of time that we have to interact with each other, and Slack doesn't really cure that," Lerman says. Software might take some of the friction out of getting work done, but it is still work.

MIGUEL PORLAN; DATA SOURCES: EMARKETER, GALLUP, RADICATI GROUP

The author could keep his hands off the wheel as his borrowed Tesla maneuvered itself through Los Angeles.

Tesla Autopilot

The electric-vehicle maker sent its cars a software update that suddenly made autonomous driving a reality.

By Ryan Bradley
Photographs by Julian Berman

Breakthrough
A car that drives itself safely in a variety of conditions.

Why It Matters
Car crashes caused by human error kill thousands of people a day worldwide.

Key Players in Autonomous Driving
- Ford Motor
- General Motors
- Google
- Nissan
- Mercedes
- Tesla Motors
- Toyota
- Uber
- Volvo

In October 2014, Elon Musk's electric-car company began rolling out sedans with a dozen ultrasonic sensors discreetly placed around both bumpers and sides. For an additional $4,250, Tesla customers could purchase a "technology package" that used the sensors, as well as a camera, a front radar, and digitally controlled brakes, to help avoid collisions—essentially allowing the car to take over and stop before crashing. But mostly, the hardware sat there, waiting, waiting, and gathering reams of data. A year later, last October 14, the company sent a software update to the 60,000 sensor-laden cars it had sold in that time. The software update was officially named Tesla Version 7.0, but its nickname—Autopilot—was what stuck.

It did in fact give drivers something similar to what airline pilots employ in flight. The car could manage its speed, steer within and even change lanes, and park itself. Some of these features, like automatic parallel parking, were already on offer from other car companies (including Mercedes, BMW, and General Motors), but the self-steering was suddenly, overnight, via a software update, a giant leap toward full autonomy.

Tesla customers, delighted, posted videos of themselves on the highway, hands free, reading the paper, sipping coffee, and even, once, riding on the roof. Some of these are, it's worth pointing out, illegal acts. Autopilot existed in a legal gray area, but it was a grand gesture toward an ever nearing future, one that will reshape not just the

Autopilot could even handle twisty Mulholland Drive, though it shut itself off in the middle of particularly tight turns.

Like many other features in the car, Autopilot can be activated or shut off from a touch screen. It also turns off with a tap on the brakes.

car and our relationship with it but the road and our entire transportation infrastructure.

Which is why I jumped at the chance to borrow a car with Autopilot for a few days and drive it—or let it drive me—around Los Angeles.

Everyone wanted to know what it felt like, the strange surrender of allowing a car to take control. The only moments that seemed like magic were when the car parked itself or changed lanes, mostly because watching a steering wheel turn all on its own was unnatural and ghostly. Other than that, I was amazed by how quickly I got used to it, how inevitable it began to feel. As a Tesla engineer told me—on condition of anonymity, because the company won't let anyone but Musk speak publicly these days—the thing that quickly becomes strange is driving a car without Autopilot. "You'll feel like the car is not doing its job," he said.

The car can't start in Autopilot; it requires a set of circumstances (good data, basically) before you can engage the setting. These include clear lane lines, a relatively constant speed, a sense of the cars around you, and a map of the area you're traveling through—roughly in that order. L.A.'s abundant highway traffic is the ideal scenario for Autopilot, not simply because of all the data it makes available to the ultrasonic sensors—which use high-frequency sound waves to identify objects up to 16 feet away—but also because humans are awful in traffic. We are bad at estimating distances to begin with, and we are constantly trying to switch lanes when the next one looks faster, causing accidents in the process. With Autopilot, I no longer had to stare at the bumper ahead of me, and I could look around to see the variety of bad decisions drivers make, stopping and starting and stopping again. Meanwhile, my car accelerated and slowed more smoothly than it ever could have with me in charge.

With its incremental approach, Tesla stands in contrast to Google and other companies that have small test fleets gathering data in hopes of someday launching fully autonomous cars. For Tesla, its customers and their partially autonomous cars are a widely distributed test fleet. The hardware required for true autonomy is already in place, so the transition can play out in software updates. Musk has said that could be technically feasible—if not legally so—within two years.

The day after I returned the Tesla, my fiancée and I were on an L.A. freeway and saw someone, speeding, cross three lanes, cutting in front of several drivers. As the traffic stopped, the car behind us came in way too fast and crashed into our bumper, which fell right off. The future, I thought, was practically here, and it couldn't arrive soon enough.


UW

Power from the Air

Internet devices powered by Wi-Fi and other telecommunications signals will make small computers and sensors more pervasive.

By Mark Harris

Breakthrough
Wireless gadgets that repurpose nearby radio signals, such as Wi-Fi, to power themselves and communicate.

Why It Matters
Freeing Internet-connected devices from the constraints of batteries and power cords will open up many new uses.

Key Players in Harvesting Radio Waves
- University of Washington
- Texas Instruments
- University of Massachusetts, Amherst

Even the smallest Internet-connected devices typically need a battery or power cord. Not for much longer. Technology that lets gadgets work and communicate using only energy harvested from nearby TV, radio, cell-phone, or Wi-Fi signals is headed toward commercialization. The University of Washington researchers who developed the technique have demonstrated Internet-connected temperature and motion sensors, and even a camera, powered that way.

Transferring power wirelessly is not a new trick. But getting a device without a conventional power source to communicate is harder, because generating radio signals is very power-intensive and the airwaves harvested from radio, TV, and other telecommunication technologies hold little energy.

Shyamnath Gollakota and his colleague Joshua Smith have proved that weak radio signals can indeed provide all an Internet gadget needs, using a principle called backscattering. Instead of generating original signals, one of their devices selectively reflects incoming radio waves to construct a new signal—a bit like an injured hiker sending an SOS message using the sun and a mirror. A gadget using the technique absorbs some energy from the signal it is modifying to power its own circuits.

"We can get communication for free," says Gollakota. RFID chips for the contactless smart cards used in mass transit also rely on backscattering, but they require specialized reader devices and can communicate only within a few inches because the reflected signals are weak and the reader itself presents interference.

One version of the University of Washington technology, dubbed passive Wi-Fi, is being commercialized through a spin-off company, Jeeva Wireless. It lets battery-free gadgets connect with conventional devices such as computers and smartphones by backscattering Wi-Fi signals. In tests, prototype passive Wi-Fi devices have beamed data as far as 100 feet and made connections through walls. Doing that requires altering the software of a Wi-Fi access point to generate an extra signal for passive Wi-Fi devices to use, very slightly increasing its power consumption.

Smith says that passive Wi-Fi consumes just 1/10,000th as much power as existing Wi-Fi chipsets. It uses a thousandth as much power as the Bluetooth LE and ZigBee communications standards used by some small connected devices and has a longer range. A device using passive Wi-Fi to communicate—for example, a security camera—could power its other circuits using energy harvested from the Wi-Fi signals it is backscattering, or by feeding on other signals such as TV and radio broadcasts.

The researchers believe that tiny passive Wi-Fi devices could be extremely cheap to make, perhaps less than a dollar. In tomorrow's smart home, security cameras, temperature sensors, and smoke alarms should never need to have their batteries changed.

ILLUSTRATION BY BRENDAN MONROE; PHOTOGRAPH BY DANIEL BERMAN

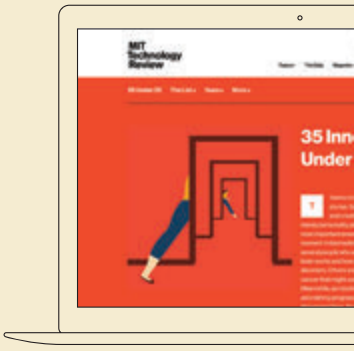

MIT Technology Review

Become an Insider and Get:

+ **1 year (6 issues)** of the magazine
+ **1 year** of online access
+ **Online story archive**: 1997-present
+ **Video**

And more

Subscribe Now

technologyreview.com/getinsider

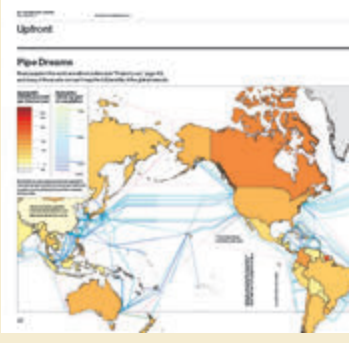

MIT Technology Review
BUSINESS REPORT

CONTENTS

Cyber Survival

With cyberattacks getting worse, the urgent need today is for faster responses, smarter technologies, and wider encryption.

SACHIN TENG

The Big Question

Cybersecurity: The Age of the Megabreach

We haven't stopped huge breaches. The focus now is on resilience, with smarter ways to detect attacks and faster ways to respond to them.

● In November 2014, an especially chilling cyberattack shook the corporate world—something that went far beyond garden-variety theft of credit card numbers from a big-box store. Hackers, having explored the internal servers of Sony Pictures Entertainment, captured internal financial reports, top executives' embarrassing e-mails, private employee health data, and even unreleased movies and scripts and dumped them on the open Web. The offenders were said by U.S. law enforcement to be working at the behest of the North Korean regime, offended by a farcical movie the company had made in which a TV producer is caught up in a scheme to kill the country's dictator.

The results showed how profoundly flat-footed this major corporation was. The hack had been going on for months without being detected. Data vital to the company's business was not encrypted. The standard defensive technologies had not worked against what was presumed to have been a "phishing" attack in which an employee clicked a link that downloaded powerful malware. Taken together, all this showed that many of today's technologies are not adequate, that attacks can now be more aggressive than ever, and that once breaches occur, they are made worse by slow responses.

The Sony hack was one in a series of recent data breaches—including many "megabreaches," in which at least 10 million records are lost—that together reveal the weakness of today's cybersecurity approaches and the widening implications for the global economy. In 2015, the U.S. Office of Personnel Management was hacked, exposing 21.5 million records, including background checks on millions of people—among them copies of 5.6 million sets of fingerprints. Later in the year, 37 million visitors to Ashley Madison, a dating site for people seeking extramarital affairs, learned that their real e-mail addresses and other data had been released. The theft of data from 83 million customers of Wall Street giant J.P. Morgan, allegedly by an Israel-based team trying to manipulate the stock market, revealed chilling possibilities for how cyberattacks could undermine the financial sector.

Since companies and other organizations can't stop attacks and are often reliant on fundamentally insecure networks and technologies, the big question for this report is how they can effectively respond to attacks and limit the damage—and adopt smarter defensive strategies in the future. New approaches and new

Cyber Breaches Hit Staggering Levels

Exceptionally harmful hacks have recently struck organizations in the global insurance, finance, telecom, and entertainment industries and at the heart of a U.S. federal agency—inflicting hundreds of millions of dollars in damage and added costs.

		How They Were Exploited	Data Stolen and Scale	Costs	Suspected Culprit
7/2014	**JPMORGAN CHASE** New York City	Two-factor authentication upgrade not fully implemented.	Names, addresses, and phone numbers of 76 million household and seven million small-business accounts.	The company says it plans to spend $250 million annually on security.	Three people have been charged with the attack as part of a stock manipulation scheme.
11/2014	**SONY PICTURES ENTERTAINMENT** Culver City, CA	Malware and lack of intrusion detection.	E-mails, salary information, and terabytes of other data, including movie scripts and contracts.	$41 million, according to public filings.	North Korean regime.
2/2015	**ANTHEM HEALTH** Indianapolis	Malware specifically designed to attack the company.	Names, birth dates, addresses, employment information, and Social Security numbers for 78 million people.	Much or all of the $100 million value of its cyberinsurance policy.	China-based hackers, suspected to be affiliated with the government.
6/2015	**U.S. OFFICE OF PERSONNEL MANAGEMENT** Washington, D.C.	Likely social-engineering attacks and lack of modern intrusion detection services.	A mix of names, birth dates, addresses, fingerprints, and background information on as many as 21.5 million people.	More than $133 million just for credit monitoring for victims.	China-based hackers, suspected to be affiliated with the government.
7/2015	**ASHLEY MADISON** Toronto	Unknown, but attackers cited weak passwords and almost nonexistent internal security.	Names, addresses, birth dates, phone numbers, and credit card history of 37 million users, plus the CEO's e-mails.	Unknown. The company faces numerous lawsuits.	A previously unknown group that calls itself Impact Team.
9/2015	**T-MOBILE US** Bellevue, Washington	Security weaknesses at a partner (Experian) that was managing credit check data.	Names, birth dates, addresses, and Social Security and driver's license numbers of 15 million people.	Experian has spent at least $20 million on credit monitoring and other corrective actions.	Unknown.
10/2015	**TALKTALK TELECOM** London	Distributed-denial-of-service attack and malicious code.	Names, birth dates, addresses, and phone numbers of more than 150,000 customers.	About $50 million in lost sales and incident response costs.	A teenager in Northern Ireland.

DATA SOURCES: SEC FILINGS, COMPANY STATEMENTS, EXPERIAN, NEW YORK TIMES, INSURANCE INSIDER

ways of thinking about cybersecurity are beginning to take hold. Organizations are getting better at detecting fraud and other attacks by using algorithms to mine historical information in real time. They are responding far more quickly, using platforms that alert security staff to what is happening and quickly help them take action. And new tools are emerging from a blossoming ecosystem of cybersecurity startups, financed by surging venture capital investment in the area.

But hindering progress everywhere is the general lack of encryption on the devices and messaging systems that hundreds of millions of people now use. Nearly three years ago, when National Security Agency contractor Edward Snowden revealed that intelligence agencies were freely availing themselves of data stored by the major Internet companies, many of those companies promised to do more to encrypt data. They started using encryption on their own corporate servers, but most users remain exposed unless they know to install and use third-party apps that encrypt their data.

$3.79 million
Average cost of a data breach

All these measures will help protect data in today's relatively insecure networks. But it's clear that the very basics of how networked technologies are built need to be rethought and security given a central role. A new national cybersecurity strategy is expected to chart an R&D plan to make sure software is verifiably secure and that users know when it's not working.

There's a big opportunity: the number of Internet-connected devices—not including smartphones, PCs, and tablets—could reach two billion in just five years. A 2015 McKinsey report predicts that this will become a multitrillion-dollar industry by 2025. All these new devices will present an opportunity to build things robustly from the start—and avoid having them play a role in Sony-like hacks in the future. *—David Talbot*

DATA SOURCE: SURVEY BY PONEMON INSTITUTE AND IBM; ILLUSTRATION: ROBERT RISKO

Expert Q&A

Why We're So Vulnerable

An expert in U.S. national cybersecurity research and policy says the next generation of technology must have security built in from the very start.

● In an age of continuing electronic breaches and rising geopolitical tensions over cyber-espionage, the White House is working on a national cybersecurity strategy that's expected in early 2016. Helping to draft that strategy is Greg Shannon. He was until recently chief scientist at Carnegie Mellon University's Software Engineering Institute and is now on leave to serve as assistant director for cybersecurity strategy at the White House Office of Science and Technology Policy.

In an interview with *MIT Technology Review* senior writer David Talbot, Shannon explained that dealing with today's frequent breaches and espionage threats—which have affected federal agencies as well as businesses and individuals—requires fundamentally new approaches to creating all kinds of software. Fixing the infrastructure for good may take two decades.

Cybersecurity has long been a serious worry. Have recent events really changed the game?

If you just consider the attack on Sony—it was a watershed event. The scale, scope, and cost were enormous. And it revealed how tightly cybersecurity and our economy are interrelated—and that the health of the economy is now potentially at stake.

Why are huge breaches like these happening? Are the billions of dollars spent on new security technologies in recent years not working?

It's more that the incentives to wage malicious cyber activities keep skyrocketing. In the early years of the Internet, the improved efficiencies from networked IT infrastructure far outweighed the security risks created by this infrastructure. Threats were always there, but it was okay to use patches. Today what's available online, and its value, keep increasing exponentially—and so do the incentives to exploit systems and steal data. What we are seeing are the results; absolutely, the threats and the attacks are bigger than they've ever been. And this hasn't been foremost in the mind-set of most companies producing software infrastructure or Internet services.

What is the underlying technology problem?

The answer might sound abstract and dry, but it has to do with efficacy and efficiency. On efficacy, how do you know that installing a new security technology is better than doing nothing? You often don't. And on efficiency, the usual approach is that you fix a newly discovered problem so the adversary doesn't use that method anymore. But at the end of the day this doesn't achieve much, because it doesn't create a general, systemic solution. It's not efficient.

We need to restructure how we build software, and develop security systems that have evidence that they actually add value. This requires rigor in how the billions of lines of code that run our networked infrastructure are actually written and updated.

The only places where software writing is truly rigorous are places like NASA—where they are building code that must work for years and from millions of miles away. They have highly formal methods and use well-controlled tools

and special engineering to make absolutely sure that the software is reliable and bug-free.

How can we make all IT infrastructure as great as the code running a Martian probe?
Many colleagues and I are devoted to this question. First, it's important to understand that there are a number of nontechnical issues that keep everyday software from being anywhere near that good. There aren't regulations or consequences that software companies experience if there are problems down the road—with the exception of certain high-priority domains like nuclear power plants or air traffic control.

The emergence of an Internet of things—interconnecting billions of devices—provides an opportunity to do things correctly from the start.

So on the policy side you need to consider incentives for everybody to write better code—it could be because of liability, regulations, or market mechanisms. And on the technology side you need to create market incentives so rigorous software development methods, like the ones NASA uses, become far more efficient and easier for everyone to use. Congress, in the 2014 Cyber Security Enhancement Act, asked for a federal cybersecurity R&D strategic plan, and that plan is being drafted, for release by early 2016.

And while it will always be true that malicious insiders or human error can create problems, great software can to a large extent deal with that, too, by creating clear access rules and sending alerts when anything anomalous happens.

Meanwhile, what can companies do to protect themselves?
Every company, from the smallest to largest, should use best practices, taking into account each company's particular assets, threats, and cybersecurity capabilities. To be sure, many systems are inherently weak. Most systems have millions of lines of code, and the typical rate for a software bug is one per 1,000 lines of code. Even if one out of a hundred of these bugs winds up creating a security vulnerability, that's a density you can't really keep up with. But if companies follow best practices, they can become much better protected—and eventually avoid more [hacks like the one on] Sony.

We aren't getting NASA-level software, but is anyone doing it right?
One simple measure that is clearly critically necessary is that products need a way to have regular and secure software updates. One can argue that companies such as Tesla and Google and Apple—and, to a large extent, Microsoft—are doing that. Google Chrome updates happen in the background; it doesn't even ask you for permission anymore.

The Apple iOS infrastructure does a good job of not requiring everyday app developers to worry about many, but not all, security issues. With Tesla, updates can happen when you charge the car.

What's the biggest opportunity right now to shape a more secure future?
The emergence of an Internet of things—interconnecting billions of devices—provides an opportunity to do things correctly from the start. Networked devices in cars and homes, and wearable devices, could introduce a multitude of new attack vectors, but if we get things right with these devices and cloud-based technologies, we can make sure the next generation of technology will have security built in.

How long until the efforts you've been talking about will make our networked infrastructure able to withstand the heightened incentives to attack it?
For the most critical components in areas like the electric grid and large industrial systems, five to 10 years is feasible. To be pervasive it will take 20 or more years.

Venture Capital

Venture Capitalists Chase Rising Cybersecurity Spending

Investors have been pouring money into companies selling "next-generation" security products.

● The rash of headline-grabbing cyberattacks on major companies over the past few years has made one thing abundantly clear: it's not enough to rely only on traditional security tools. To venture capitalists, that means there's money to be made by betting on startups developing new ones.

VCs are hoping to get a piece of companies' increased spending on cybersecurity. In 2014 Gregg Steinhafel, the CEO of Target, became the first head of a major company to lose his job over a data breach. Now, worried company leaders are giving their security units a "blank check," says Scott Weiss, a general partner who specializes in security at the venture capital firm Andreessen Horowitz: "The CEO has said, 'Look, whatever you need, you've got.'"

$3.3 billion
2015 VC investment in cybersecurity

Today's advanced threats are much too sophisticated for traditional tools like antivirus software and firewalls. Not wanting to buy obsolete products, security executives are increasingly venturing into agreements with cybersecurity startups. To Weiss and other venture investors, that kind of customer demand is an investment opportunity. According to CB Insights, the global VC community

Rising Venture Capital Interest in Cybersecurity Startups

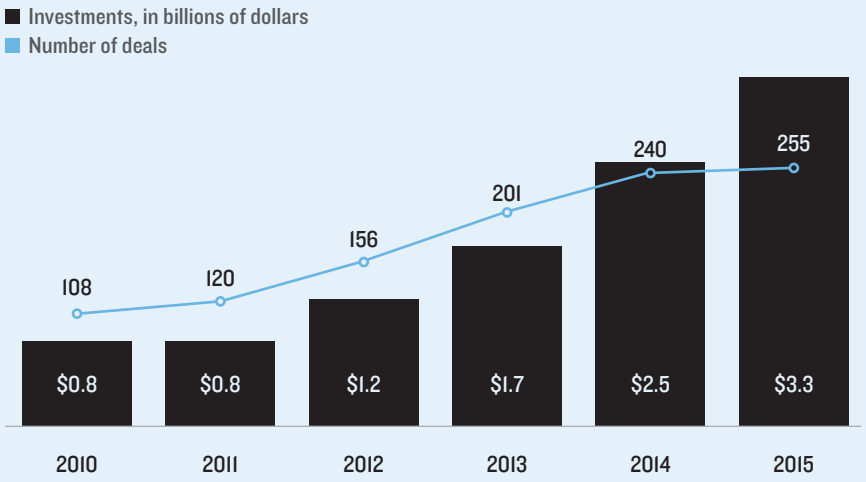

poured a record $2.5 billion into cybersecurity companies in 2014, a strong year for IT startups in general and software in particular. Security companies raised another $3.3 billion in 2015.

The problems these startups are trying to solve are complex. The bad guys do have better weapons, but business systems are also becoming vulnerable in new ways. Businesses are relying more on cloud services and connecting more "things" to the Internet, and their employees are using more connected devices.

Before a few years ago, the conventional approach to security entailed basically building a wall around valuable data and using software to detect known signatures of malicious code. Then security researchers began finding extremely complex malware, derived from government-designed exploits and sophisticated enough to circumvent traditional antivirus tools. This new generation of malware can be custom-built for a specific network and more precisely controlled by its human operators.

Dealing with such specialized, fast-evolving adversaries requires changing the security paradigm from prevention to "active cyberdefense," says Nicole Eagan, CEO of Darktrace, a two-year-old company based in Cambridge, United Kingdom. Hackers are going to get in, so the trick now is to find them "in near real time as they are moving subtly and silently around your network" and catch them before they do any real damage, she says.

Many of the new startups are focused on trying to detect hackers in real time as they enter and move subtly through a corporate network.

A number of companies, taking a range of different approaches, promise that their detection technologies can do this. Darktrace, which has raised $50 million in VC funding, relies on advanced machine-learning technology to analyze raw network traffic and, as Eagan explains, "determine a baseline for what's normal" for every person using the network so that it can detect abnormal behavior.

Not only are the threats more numerous and advanced, but companies must also secure networks that are growing more complex and massive. Every device on a network is a potential target for hackers, and new security technologies focused on them are getting lots of attention from investors.

A company called Tanium, which is now valued at $3.5 billion after its most recent round of VC investment, has technology that allows network operators to ask questions about what's happening in any one of millions of devices on a network. They get an answer within 15 seconds and can quickly take action—for example, by quarantining an infected computer.

Security investors are also focused on the fact that businesses and organizations are putting more and more data in the cloud. In response to this trend, a new breed of cloud security companies are offering services such as novel encryption schemes and technologies for continuously monitoring what goes on in a company's cloud servers.

With so much funding available, the burgeoning cybersecurity startup scene is chaotic. Greg Dracon, a partner at .406 Ventures who has invested in several security companies, thinks a consolidation cycle may already be starting. Bigger companies are buying up individual technologies and could eventually offer suites of products, he says.

Dracon thinks all this investor attention is driving prices too high overall, and that the market has gotten ahead of itself, at least for the near term. However, the security market itself has another decade of growth at least, he believes. "The problem set is outpacing the solution set," he says, "and I don't think there's any end in sight to that." —*Mike Orcutt*

DATA SOURCE: CB INSIGHTS

Case Study

How PayPal Boosts Security with Artificial Intelligence

The payments giant keeps fraud losses below industry averages by teaching computers to play detective.

● To PayPal, the transactions signal fraud: a U.S. user's account is accessed in the U.K., China, and elsewhere around the world. But PayPal's security

system—thanks to a growing reliance on an artificial-intelligence technology known as deep learning—is now able to spot possible fraud without making mistakes. That's because algorithms mine data from the customer's purchasing history—in addition to reviewing patterns of likely fraud stored in its databases—and can tell whether, for example, the suspect transactions were innocent actions of a globe-hopping pilot.

From a cybersecurity perspective, PayPal has a target on its back: it processed $235 billion in payments last year from four billion transactions by its more than 170 million customers. Fraud is always possible via theft of consumer data in breaches such as "phishing" e-mails that con users into entering their credentials. To keep ahead, PayPal relies on intensive, real-time analysis of transactions.

When a pattern is revealed—for example, if sudden strings of many small purchases at convenience stores turn out to be fraud—it's turned into a "feature," or a rule that can be applied in real time to stop purchases that fit this profile. "We now process thousands of 'features' in our system, compared to hundreds when the system was first put to use in 2013," says Hui Wang, the company's senior director of global risk sciences.

As a result, PayPal can now do things like tell the difference between friends buying concert tickets together and a thief making similar purchases with a list of stolen accounts. And it's all done in-house to avoid even the tiny latency that would occur if the company relied on a cloud provider. "Thousands of 'features' searching through 16 years of users' history all needs to be done in less than a second," Wang says.

As a transaction is being made, PayPal's deep-learning algorithms can search 16 years of user purchasing history and thousands of fraud patterns to spot theft while avoiding error.

Deep learning and other artificial-intelligence approaches are quickly becoming the only way to keep up with threats, she adds. They've worked to help keep PayPal's fraud rate remarkably low, at 0.32 percent of revenue—a figure far better than the 1.32 percent average that merchants see, according to a study by LexisNexis. The most recent Federal Reserve Payments Study found that $6.1 billion in fraudulent purchases were made in 2012, and the problem appears to be getting worse.

PayPal isn't the only company using deep learning to improve cybersecurity. The Israeli startup Deep Instinct has employed the technique to spot malware, claiming that this works 20 percent better than traditional approaches. And Ashar Aziz, vice chairman and founder of the security firm FireEye, said that his company has been using deep learning for everything from detecting network intrusions to rooting out phishing attacks.

Companies can further improve cybersecurity if they share data repositories on cyberattacks and fraud, says Aziz. "If you continue to get more data—and more power to process it—then you can get even better," he says. —*Michael Morisy*

Rising Crime

Fraud losses are increasing. But PayPal's rate is only 0.32%.

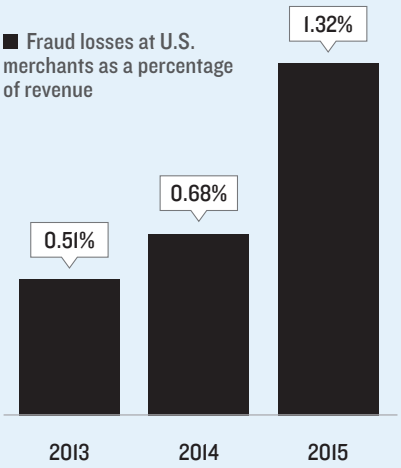

DATA SOURCE: LEXISNEXIS TRUE COST OF FRAUD REPORT

Encryption

Half-Measures on Encryption Since Snowden

Amid a wave of corporate privacy and security pronouncements, 2014 was supposed to be the "year of encryption." It didn't pan out that way.

● When the NSA subcontractor Edward Snowden released classified documents in June 2013 baring the U.S. intelligence community's global surveillance programs, it revealed the lax attention to privacy and data security at major Internet companies like Apple, Google, Yahoo, and Microsoft. Warrantless surveillance was possible because data was unencrypted as it flowed between internal company data centers and service providers.

The revelations damaged technology companies' relationships with businesses and consumers. Various estimates pegged the impact at between $35 billion and $180 billion as foreign business customers canceled service contracts with U.S. cloud computing companies in favor of foreign competitors, and as the companies poured money into PR campaigns to reassure their remaining customers.

There was a silver lining: the revelations catalyzed a movement among technology companies to use encryption to protect users' data from spying and theft. But the results have been mixed. Major service providers including Google, Yahoo, and Microsoft—who are among the largest providers of cloud- and Web-based services like e-mail, search, storage, and messaging—have indeed encrypted user data flowing across their internal infrastructure. But the same isn't true in other contexts, such as when data is stored on smartphones or moving across networks in hugely popular messaging apps like Skype and Google Hangouts. Apple is leading the pack: it encrypts data by default on iPhones and other devices

running newer versions of its operating system, and it encrypts communications data so that only the sender and receiver have access to it.

But Apple products aren't widely used in the poor world. Of the 3.4 billion smartphones in use worldwide, more than 80 percent run Google's Android operating system. Many are low-end phones with less built-in protection than iPhones. This has produced a "digital security divide," says Chris Soghoian, principal technologist at the American Civil Liberties Union. "The phone used by the rich is encrypted by default and cannot be surveilled, and the phone used by most people in the global south and the poor and disadvantaged in America can be surveilled," he said at *MIT Technology Review*'s EmTech conference in November.

Pronouncements on new encryption plans quickly followed the Snowden revelations. In November 2013, Yahoo announced that it intended to encrypt data flowing between its data centers and said it would also encrypt traffic moving between a user's device and its servers (as signaled by the address prefix HTTPS). Microsoft announced in November and December 2013 that it would expand encryption to many of its major products and services, meaning data would be encrypted in transit and on Microsoft's servers. Google announced in March 2014 that connections to Gmail would use HTTPS and that it would encrypt e-mails sent to other providers who can also support encryption, such as Yahoo. Finally, in 2014, Apple implemented the most dramatic change of all, announcing that the latest version of iOS, the operating system that runs on all iPhones and iPads, would encrypt user data (including content on Apple's Messages app) with the user's chosen passcode, not with a key the device could access. This made it far more difficult for a hacker, an insider at Apple, or even government officials with a court order to gain access to user data. Apple had provided end-to-end encrypted video and text messaging since 2010 and 2011.

Google, Microsoft, and Yahoo don't provide such robust methods to encrypt users' data. But users can turn to a rising crop of free third-party apps, like ChatSecure and Signal, that support such encryption and open their source code for review. Relatively few users take the extra step to learn about and use these tools. Still, secure messaging apps may play a key role in making it easier to implement wider encryption across the Internet, says Stephen Farrell, a computer scientist at Trinity College Dublin and a leader of security efforts at the Internet Engineering Task Force, which develops fundamental Internet protocols. "Large messaging providers need to get experience with deployment of end-to-end secure messaging and then return to the standards process with that experience," he says. "That is what will be needed to really address the Internet-scale messaging security problem." —*David O'Brien*

Sobering Message

Unlike other big players, Apple does a thorough job of encrypting messages.

	Encrypted in transit	Encrypted end-to-end	Supports verification of contacts' identities	Deletes encryption keys after use	Code open to independent review	Properly documents security	Recently audited
iMESSAGE (APPLE)	●	●		●		●	●
FACETIME (APPLE)	●	●		●		●	●
HANGOUTS/CHAT (GOOGLE)	●						●
SKYPE (MICROSOFT)	●						
YAHOO MESSENGER	●						

ILLUSTRATION: PATRICK KYLE; DATA SOURCE: ELECTRONIC FRONTIER FOUNDATION

Incident Response

New Rapid Response Systems Blunt Cyberattacks

Limiting damage from attacks requires far faster reactions, quick notification of victims, and adherence to regulations. Managing all that can be tricky.

● One reason breaches do so much damage is that they often remain undiscovered for months—an average of more than 200 days, according to research by the security firm Mandiant. Over time, a once-minor breach can become a catastrophe.

Sometimes the intrusion is hard to spot because the hacker has stolen legitimate credentials. Other times subtle hints of unusual network activity that might have revealed the attack are simply missed.

When such clues go unnoticed, it is often because large corporate security systems are so complex. It's not unusual for a big company to use 70 or more different security monitoring tools made by many different companies and adopted over time—intrusion detectors, firewalls, Web-page monitors, spam filters, and many others. This common situation "is a huge problem," says Jon Oltsik, cofounder of Enterprise Strategy Group, an IT research firm. "It depends on very, very smart people to figure out what each system is telling them and put together the total picture."

One solution is for companies to replace whatever they've already installed with integrated systems from giant vendors like IBM, Cisco, and Raytheon. But that can be expensive and impractical for many.

So a growing crop of startups and research projects are beginning to offer approaches aimed at making it easier to tie existing systems together, while also making it possible to respond to attacks quickly and appropriately.

One early entrant, Resilient Systems, a startup in Cambridge, Massachusetts, captures data from a variety of sources and provides a single dashboard displaying all warnings. Then it presents a checklist of actions that must be taken, both to fix the problem and also to do things like notify the U.S. Federal Trade Commission or comply with state-by-state consumer notification laws. *—David Talbot*

Incident Response: An Emerging Field

Startup companies are joining major federal research efforts and open-source technologies to create quick and effective responses to cyberattacks.

	Technology	Date
RESILIENT SYSTEMS Cambridge, MA	Cyberattack detection and response platform	Launched in 2011
INVOTAS Alexandria, VA	Cyberattack detection and response platform	Launched in 2014
HEXADITE Tel Aviv, Israel	Cyberattack detection and response platform	Launched in 2014
NETFLIX Los Gatos, CA	Open-source software for cyberattack detection and response	Released in 2015
U.S. DEPARTMENT OF DEFENSE/JOHNS HOPKINS UNIVERSITY	Adaptable cyber response systems	Ongoing federal research project
U.S. DEPARTMENT OF HOMELAND SECURITY	Automated cybersecurity for businesses	Recent government/industry collaboration

Emerging Technologies

Finding Insecurity in the Internet of Things

The world of connected devices is growing fast, but how secure is it?

● As we connect everything from Barbie dolls to front-door locks and cars to the Internet, we're creating more—and possibly more dangerous—potential ways for cyberattackers to wreak havoc.

Security researchers have reported on the ease with which you can break into a range of connected gadgets like baby monitors and cars. This past summer a piece in *Wired* showed how a software bug could be exploited to control a Jeep driving down the highway. (Jeep owner Chrysler quickly fixed the bug.)

Mika Ståhlberg, director of strategic threat research at the Finnish security company F-Secure, points out that while a hacked credit card may be a headache, a hacked smart lock could open your home to burglars.

A number of startups have started offering security for the Internet of things. In November, F-Secure announced a product called Sense that can monitor Internet-connected devices like smartphones, smart lights, and baby monitors. The device, which should be available in the spring, keeps an eye on network metadata—which includes information like where data is going or coming from, and how much is being sent overall—and blocks activity thought to be malicious. Atlanta-based Bastille, meanwhile, uses sensors to keep track of connected devices by measuring the electromagnetic signatures of different devices in an office. The sensors can track devices that use communication protocols like Wi-Fi and low-energy Bluetooth or work over cellular networks, and its software can tell

DATA SOURCE: SYMANTEC

where they are to within three meters. Bastille's tactic of scanning a wide spectrum of radio frequencies suits Internet-connected gadgets since they are designed using many different protocols.

The potential range of attack targets is rising: Gartner, the market research firm, predicts that by 2020, almost 21 billion gadgets will be connected to the Internet, up from 4.9 billion today. "This is the World Wide Web of 1994, 1995. We know it's going to be big," says Phil Levis, an associate professor at Stanford who co-directs the university's Secure Internet of Things Project. "It's going to be a security train wreck, much as the Web was for 10 years or so until people figured it out."

Levis isn't convinced monitoring is the best approach, because behavior variations will only show up after a device has been compromised or an attack has occurred, he says. What really needs to happen, he says, is for device manufacturers to write secure software in the first place. The Internet is in some ways more secure now than two decades ago, because developers are more careful and clean up dangerous code. These lessons have yet to be picked up by many Internet-of-things developers, he says. *—Rachel Metz*

21 billion

Number of Internet-connected "things" expected in 2020

China

China Hit by Rise of Attacks

China sees a major increase in infections on file-sharing sites and more targeted, localized malware threats.

● China-based hackers are sometimes accused of being behind major external attacks like the one on the U.S. Office of Personnel Management, as well as acts of corporate espionage. But China has worsening internal problems, too.

In September, a counterfeit copy of Apple's Xcode software development tool was offered on a local file-sharing site, leading to infections on iPhone apps created with the fake tool. The hack, which ended up affecting more than 100 million mostly China-based iPhone users, was Apple's biggest security breach to date.

A possibly even larger hack was an October attack on NetEase, one of the top social-media and news platforms in China. A hack of its 163.com e-mail system, which is still under investigation, potentially exposed the aliases, security questions and answers, passwords, and other data of hundreds of millions of primarily Chinese users.

Hong Jia, a cofounder of the China-based threat intelligence firm ThreatBook and former cybersecurity expert at Microsoft, says companies and individuals in China are beginning to wake up to the threat. "Enterprises [in China] know that someday they will get targeted and a whole company can be exposed by an attack," Hong said in an interview at the Association of Anti-Virus Asia Researchers International Conference, held in December in Danang, Vietnam.

According to a survey by auditing firm PricewaterhouseCoopers, over the past year companies in China and Hong Kong saw around 1,245 attacks each on average, compared with 241 the year before. In addition to big hacks like the iPhone incident, Chinese companies have experienced a rapidly rising number of attacks that use so-called social engineering to trick individuals into clicking links that download malware onto the user's computer. "The threats you see in China are really, really targeted," Ingvar Froiland, director and general manager for the security company F-Secure, said in an interview at the Danang conference. Froiland said the threats are often language-specific or event-specific—such as targeted attacks during Chinese New Year and other holidays. He added that they also may be system- and application-specific: for example, they are sometimes launched through games that may not be used widely outside China, or through file-sharing sites accessed mainly by Chinese users.

Chinese authorities even discovered a "hacking village" last year. In a mostly rural area bordering Vietnam, large numbers of people were involved in cybercrime, cyberfraud, and hacking, often using the popular QQ instant messaging software run by Tencent, one of the world's biggest Internet companies.

At the Danang conference, Liu Zhao, an antimalware analyst at Tencent, said he has been finding increasing numbers of new tricks deployed by hackers in China, including malicious files masquerading as harmless icons attached to documents sent to specific victims. Real-world parent-teacher, school-student, or business-consumer relationships—often discovered from stolen e-mails—are sometimes used for extortion, he added.

To fight targeted attacks, Hong said, analysts are working on analyzing traffic flowing from computer addresses and domain names to find the source, such as the hacking village. "We can see ... what person might be behind it," Hong said. Adding to China's woes is that citizens often do not add protections to their mobile devices. Worldwide, "awareness of threats to mobile devices is not there yet," Froiland said. *—Michael Standaert*

ACCESS THE FULL REPORT ONLINE

Why is it so difficult to estimate how much cybercrime costs us each year? Why hasn't a robust market for cyberinsurance developed yet? Many such questions face cybersecurity in our age of the megabreach.

Europe/U.S. Data Friction

Tallying the Cost of Cybercrime

Dangerous Back Doors

Insuring the Internet of Things

Upcoming Events, Readings, and More

technologyreview.com/business

Reviews

Should Silicon Valley Go to War?

Politicians are trying to recruit technology companies to help fight ISIS. Does it make sense?

By Fred Kaplan

On Friday, January 8, several high-level officials from the Obama administration—including the attorney general, the White House chief of staff, and the directors of the FBI and the NSA—met at a federal office in San Jose with senior executives from Facebook, Twitter, Microsoft, LinkedIn, YouTube, and Apple (including CEO Tim Cook himself). On the agenda for the discussion, according to a one-page memo widely leaked to the press, was this question: "How can we make it harder for terrorists to [use] the Internet to recruit, radicalize, and mobilize followers to violence?"

For the previous month, since the ISIS-inspired shootings in San Bernardino, California, President Obama, as well as some of the candidates vying to succeed him, had been calling on Silicon Valley to join the government in this fight. As Hillary Clinton put it in a campaign speech, "We need to put 'the great disrupters' at work disrupting ISIS." In one of the Republican presidential debates, Donald Trump said he would ask "our brilliant people from Silicon Valley" to keep ISIS from using the Internet—a notion that reflected a misunderstanding of how the Internet works but also a widespread desperation for Silicon Valley to do *something*.

But what do Obama, Clinton, Trump, and the other politicians have in mind? How would the executives respond, and how *should* they respond? Many tech entrepreneurs—libertarian in leanings and especially leery of open collusion with Washington since Edward Snowden's revelations—question whether government has any business putting private industry to work on such a venture, which could rub up against the First, Fourth, and Fifth Amendments. And if some coöperative strategy could be mounted, quite apart from any philosophical considerations, would it have much effect?

Hillary Clinton: "National Security and the Islamic State"
November 19, 2015, at the Council on Foreign Relations

Donald Trump: Republican debate
December 15, 2015

President Obama: "Keeping the American People Safe"
December 6, 2015

The main thing, Clinton emphasized, was that, above and beyond questions about specific plans, "the tech community and the government have to stop seeing each other as adversaries."

This enmity, especially from the techies toward the spies, is a fairly new phenomenon. Telecommunications companies have a history of coöperating with U.S. intelligence agencies that dates back to the 1920s, when the Cipher Bureau, which grew out of a World War I espionage unit, persuaded Western Union to grant its agents access to all telegrams and telegraphs. Starting in the 1950s, with the founding of the National Security Agency, AT&T and later the Baby Bells allowed signal-intelligence crews to tap into phone lines. A whole industry grew up to build listening posts, dishes, and satellites that intercepted radio and microwave signals. When the world went digital, the new Internet and cellular companies continued the tradition of complicity—sometimes under court order, more often willingly. Favors were reciprocated. For instance, two senior NSA officials told me that when Microsoft released its first Windows software, the agency's Information Assurance Directorate inspected the product (as it was obligated to do

GUY BILLOUT

before approving it for procurement by the Defense Department), found 1,500 points of vulnerability, and helped patch almost all of them (leaving a few of the gaps open so the NSA could exploit them in adversaries' computer systems).

The Snowden leaks, in June 2013, exposed the extent of this arrangement, embarrassing several executives and fomenting fears that consumers abroad might shop elsewhere because they'd assume that American-made products had built-in back doors for NSA intruders. Apple declared its independence in particularly dramatic fashion, designing the encryption for its IOS 8 operating system, released in 2014, in a way that let consumers set their own passcode: Apple couldn't hand the NSA a key, because it didn't *have* the key.

The rise of ISIS has altered this atmosphere and softened the hostility. Even the most freewheeling executives in Silicon Valley have no desire to let outfits like ISIS use their networks, sites, or servers at will. So there is, in principle, a revived willingness to coöperate with Washington—or at least, for the moment, to engage in dialogue—on a cyber counterterrorist campaign.

Some coöperation is going on already. Facebook and Twitter have taken steps to spot terrorist posts and take them down, though their efforts have proved futile: new sites and pages spring up as fast as the old ones are shuttered. But there are other ways to disrupt nefarious plots online; and though they weren't discussed in detail at the San Jose meeting, the history of what has variously been called "information operations," "information warfare," and "cyberwarfare" suggests a wide range of possible techniques.

In 2007, four years into the Iraq War, U.S. forces started making headway: American casualties plummeted, insurgent casualties soared. The official story credited the turnaround to President George W. Bush's troop surge and General David Petraeus's adoption of a counterinsurgency strategy. There's something to that story, but another factor, which several officials told me about but no one discussed openly, was a cyberwar campaign.

"How can we help others to create, publish, and amplify alternative content that would undercut [ISIS]?"

U.S. Special Forces captured insurgent computers. NSA analysts, deployed on the ground, downloaded insurgents' usernames and passwords, then sent phony e-mails to insurgent fighters, ordering them to meet at a certain location at a certain time—where members of the Special Forces would be waiting to kill them. In the course of several months, 4,000 insurgents were killed in this fashion. (So were 22 NSA analysts, mainly by roadside bombs as they accompanied troops on missions to capture computers.)

The sort of counter-ISIS program discussed by senior officials and Internet executives wouldn't go quite that far. Actually killing jihadists in this manner would require troops on the ground and (like all cyber-offensive activities that involve killing people or destroying objects) presidential authorization. But it would not be a stretch (and would require no permission from political higher-ups) to capture—or hack into—ISIS computers, track the Twitter feeds and Facebook pages involved in recruiting new fighters, and follow the ensuing e-mails to and from those who respond. The resulting information could be gathered strictly as intelligence—to analyze the personality types, or identify the specific individuals, who are lured. Or messages between the recruiter and the recruited could be disrupted or distorted in a way that undermined the movement's appeal. Or the pages could be flooded with comments by Muslims—real or invented—disputing or ridiculing the recruiter's message, snapping susceptible readers out of their reverie or making them think twice before booking a plane ride and taking up arms.

Matt Devost, a cybersecurity specialist who ran the Terrorism Research Center for 13 years, says, "Studies of group dynamics indicate that if dissenting voices are introduced, they can diminish the appeal of propaganda."

The Obama administration may at least be considering these sorts of approaches. The one-page agenda for the January 8 meeting in San Jose asked: "In what ways can we use technology to help disrupt paths to radicalization to violence, identify recruitment patterns, and provide metrics to help measure our efforts?" And: "How can we help others to create, publish, and amplify alternative content that would undercut [ISIS]?"

To some extent, some pushback is already happening spontaneously. In 2014, a message by Abu Bakr al-Baghdadi, the self-proclaimed caliph of the Islamic State, was tweeted: "We urgently call upon every Muslim to join the fight, especially those in the land of the two shrines" (by which he meant Saudi Arabia). Someone named Mohsin Arain replied, "Sorry mate, I don't want to risk dying before the next Star Wars comes out." Another, Zay Zadeh, posted, "Sorry … I'm busy being a real Muslim, giving to charity, etc. Also, your dental plan sucks." Still another, Hossein Aoulad, responded, "Mum just made couscous, next time maybe." Imagine if hundreds of counter-messages flooded an ISIS message board, and if at least some of them were designed to appeal to the sorts of people whom intelligence analysts had pegged as susceptible to recruitment. Keeping a propaganda line open—in order to track and possibly manipulate its contents and controllers—

might be far more effective than a whack-a-mole attempt to shut it down.

Another virtue of this approach is that even if the jihadist leaders suspected some of the dissenting voices weren't genuine, even if they knew the West was using their sites to mount a counter-propaganda campaign, there wouldn't be much they could do about it; their anonymous readers, in bedrooms and basements around the world, would regard them as real.

Who would decide to run this sort of campaign—the government or the Internet companies? Law enforcement and intelligence agencies have the necessary resources, personnel, and institutional mandate. But the companies would need to play a role as well: they own the networks. Their role could be passive—for instance, receiving notice that some agency is monitoring or disrupting a particular site, so they don't shut it down. Or it could be active, ranging from providing new ideas (their business model encourages innovative thinking much more than government bureaucracies do) to carving out a back door in the architecture of a site, a server, or a network so that a spy agency's hackers could enter. Whatever the precise arrangement, the government needs Silicon Valley to at least be a partner. In that sense, Hillary Clinton got it right when she called on each side to stop seeing the other as an adversary.

This dialogue is in an early phase. The January meeting in San Jose, according to one attending official who was not authorized to speak on the record, amounted to a "preliminary discussion," which was conducted on "an unclassified basis"—meaning none of the ideas or scenarios cited above would have been outlined, except perhaps on an abstract level. But officials hope—while some libertarians fear—that the meeting may presage a softening of Silicon Valley's resistance.

Just as telecom executives in the last half of the 20th century felt moved by appeals to national security during the Cold War, so Internet executives today—after two decades of relative peace, a go-go economy, and the motto "information wants to be free"—might be lured back to pledges of allegiance, at least to some extent, by the threat of global terrorism.

Fred Kaplan is Slate*'s "War Stories" columnist and the author of* Dark Territory: The Secret History of Cyber War*, out on March 1.*

Apprentice Work

What is the potential of machine art, and can it truly be described as creative or imaginative?

By Martin Gayford

The Painting Fool

AARON

Google's Inceptionist photographs

In July 2013, an up-and-coming artist had an exhibition at the Galerie Oberkampf in Paris. It lasted for a week, was attended by the public, received press coverage, and featured works produced over a number of years, including some created on the spot in the gallery. Altogether, it was a fairly typical art-world event. The only unusual feature was that the artist in question was a computer program known as "The Painting Fool."

Even that was not such a novelty. Art made with the aid of artificial intelligence has been with us for a surprisingly long time. Since 1973, Harold Cohen—a painter, a professor at the University of California, San Diego, and a onetime representative of Britain at the Venice Biennale—has been collaborating with a program called AARON. AARON has been able to make pictures autonomously for decades; even in the late 1980s Cohen was able to joke that he was the only artist who would ever be able to have a posthumous exhibition of new works created entirely after his own death.

The unresolved questions about machine art are, first, what its potential is and, second, whether—irrespective of the quality of the work produced—it can truly be described as "creative" or "imaginative." These are problems, profound and fascinating, that take us deep into the mysteries of human art-making.

The Painting Fool is the brainchild of Simon Colton, a professor of computational creativity at Goldsmiths College, London, who has suggested that if programs are to count as creative, they'll have to pass something different from the Turing test. He suggests that rather than simply being able to converse in a convincingly human manner, as Turing proposed, an artificially intelligent artist would have to behave in ways that were "skillful," "appreciative," and "imaginative."

Thus far, the Painting Fool—described as "an aspiring painter" on its website—has made progress on all three fronts. By "appreciative," Colton means responsive to emotions. An early work consisted of a mosaic of images in a medium resembling watercolor. The program scanned an article in the *Guardian* on the war in Afghanistan, extracted keywords such as "NATO," "troops," and "British," and found images connected with them. Then it put these together to make a composite image reflecting the "content and mood" of the newspaper article.

The software had been designed to duplicate various painting and drawing media, to select the appropriate one, and also to evaluate the results. "This is a miserable failure," it commented about one effort. A skeptic might doubt whether this and other statements are anything more than skillful digital ventriloquism. But the writing of poetry is mentioned on the website as a current project—so the Painting Fool apparently aspires to be an author as well as a painter.

In the Paris exhibition, the sitters for portraits faced not a human artist but a laptop, on whose screen the "painting" took place. The Painting Fool executed pictures of visitors in different moods, responding to emotional keywords derived from 10 articles read—once again—in the *Guardian*. If the tally of negativity was too great (always a danger with news coverage), Colton programmed the software to enter a state of despondency in which it refused to paint at all, a virtual equivalent of the artistic temperament.

Arguably, the images unveiled in June 2015 by Google's Brain AI research team also display at least one aspect of human imagination: the ability to see

COURTESY OF HAROLD COHEN

One of AARON's compositions.

one thing as something else. After some training in identifying objects from visual clues, and being fed photographs of skies and random-shaped stuff, the program began generating digital images suggesting the combined imaginations of Walt Disney and Pieter Bruegel the Elder, including a hybrid "Pig-Snail," "Camel-Bird" and "Dog-Fish."

Here is a digital equivalent of the mental phenomenon to which Mark Antony referred in Shakespeare's *Antony and Cleopatra*: "Sometime we see a cloud that's dragonish/A vapour sometime like a bear or lion."

Leonardo da Vinci recommended gazing at stains on a wall or similar random marks as a stimulus to creative fantasy. There, an artist trying to "invent some scene" would find the swirling warriors of a battle or a landscape with "mountains, rivers, rocks, trees, great plains, valleys and hills." This capacity might have been one of the triggers for prehistoric cave art. Quite often a painting or rock engraving seems to use a natural feature—a pebble in the wall that looks like an eye, for example. Perhaps the Cro-Magnon artist first discerned a lion or a bison in random marks, then made that resemblance clearer with paint or incised line.

Come to that, all representational pictures—not only paintings and drawings but also photographs—depend on a capacity to see one thing, shapes on a flat surface, as something else: something in the three-dimensional world. The artificial-intelligence systems developed by the Google team are good at that. The images were created using an artificial neural net, software that emulates the way layers of neurons in the brain process information. The software is trained, through analyzing millions of examples, to recognize objects in photos: a dumbbell, a dog, or a dragon.

Collage by the Painting Fool, inspired by news from Afghanistan.

The Google researchers discovered they could turn such systems into artists by doing something like what Leonardo suggested. The neural net is provided with an image made up of a blizzard of blotches and spots, and is asked to tweak the image to bring out any faint resemblance it detects in the noise to objects that the software has been trained to recognize. A sea of noise can become a tangle of ants or starfish. The technique can also be applied to photos, populating blue skies with ghostly dogs or reworking images in stylized strokes.

The software was just as adept as Mark Antony at discerning animals in clouds. The Google team dubbed the resulting artistic idiom "Inceptionism," because the research project into neural-network architecture was code-named "Inception"—a reference to a 2010 movie of the same name about a man who penetrates deeper and deeper layers of other people's dreams. Art-historically, you might classify Inceptionism as a variant of Surrealism. René Magritte, Salvador Dalí, and Max Ernst produced numerous works of a similar type—painting a sky of musical instruments or baguettes, for instance, instead of cumulonimbus.

How good, really, is Inceptionism? Some of the pictures are striking and can be perceived in various ways—including an emphatic linear mode vaguely reminiscent of the style of Van Gogh. In some cases, they are disturbing, suggesting the kind of hallucinations described by those suffer-

COURTESY OF SIMON COLTON

ing from bad trips or the DTs: a sky filled with cycling dogs, for example, or swirling architecture covered in peering eyes.

But Inceptionist works, so far, have been too kitschy and too evidently photo-based—to my taste, anyway—to give much competition to Dalí or Magritte. Nor have the Painting Fool or most similar programs yet progressed beyond a high school or amateur-art-club level of achievement. What about the *potential* of computer art? Can artificial intelligence add to the visual arts (or, for that matter, to music and other idioms at which computers are also already adept)?

Simon Colton is conscious of the criticism—a standard one aimed at computer art—that the works of the Painting Fool are actually creations of his own. We wouldn't, he has pointed out, give the credit for a human painter's work to that artist's teacher. To which the answer is, that might depend on how far the pupil was following the teacher's instructions. Generally, the credit for a painting from a Renaissance workshop goes to the master, not to the apprentices who may have done much of the work. But in the case of Verrocchio's *Baptism of Christ* (c. 1475), we recognize the achievement of the workshop member Leonardo da Vinci, because the parts he painted—an angel and some landscape—are visibly different from the master's work. Art historians therefore classify the picture as a joint effort.

In 17th-century Antwerp, similarly, Rubens had a small factory of highly trained assistants who to a greater or lesser extent painted most of his large-scale works. The normal procedure was that the master produced a small sketch that was then blown up, under his supervision, to the size of a ceiling or an altarpiece. Some scholars believe, however, that on occasion the studio turned out a "Rubens" when the great man never even provided an initial model.

Google's neural networks hallucinate patterns and objects.

Here, the example of AARON is intriguing. Are the pictures the evolving program has made over the last four decades really works by Harold Cohen, or independent creations by AARON itself, or perhaps collaborations between the two? It is a delicate problem. AARON has never moved far out of the general stylistic idiom in which Cohen himself worked in the 1960s, when he was a successful exponent of color field abstraction. Clearly, AARON is his pupil in that respect.

One aspect of Cohen's earlier work was crucial to his taking an interest in artificial intelligence. He felt that "making art didn't have to require ongoing, minute-by-minute decision making ... that it should be possible to devise a set of rules and then, almost without thinking, make the painting by following the rule."

This approach is characteristic of a certain type of artist. The classic abstractions of Piet Mondrian from the 1920s and 1930s are a case in point. These were made according to a set of self-imposed regulations: only straight lines were allowed, which could meet only at right angles and could be depicted only in a palette of red, blue, and yellow (plus black and white).

In a rare example of an art-historical experiment, the late art critic Tom Lubbock attempted to paint some Mondrians himself by following this recipe. He duly

GOOGLE

produced several abstractions that looked quite like Mondrian's works, just not so good. The conclusion appeared to be that Mondrian was adding extra qualities—perhaps subtleties of visual balance and weighting of color—that weren't formulated in the rules.

It is unusual for art critics to try anything as practical as Lubbock's research. But lots of other people do the same kind of thing: they are called forgers, copyists, and pupils. A great deal of art consists, and always has consisted, of imitations of other work: pictures done in the manner of Mondrian, Monet, or some other great originator. Art historians spend their lives classifying artists into "circle of Botticelli," "follower of Caravaggio," etc. Already, it is clear that machines can work on this level: they can produce derivative art (which is all that 99.9 percent of human artists do). But can they do more than that?

Understandably, Cohen has thought a great deal about this. In a lecture from 2010, he posed it the other way around. Wasn't it obvious that AARON is creative? After all, he went on, "with no further input from me, it can generate unlimited numbers of images, it's a much better colorist than I ever was myself, and it typically does it all while I'm tucked up in bed." What, in fact, he asked, was his own contribution? "Well, of course, I wrote the program. It isn't quite right to say that the program simply follows the rules I gave it. The program *is* the rules."

In a way, then, AARON is functioning like a Renaissance or Baroque studio. Under Cohen's direction, it has developed to the point where it is equivalent to Rubens's studio in autonomous mode—and perhaps more. In the early years, AARON was confined to drawing outlines; Cohen then selected and sometimes added color by hand. In the '80s, Cohen began to teach it to work in color. Eventually, he developed a series of rules to enable it to compose coloristic harmonies, but he found this unsatisfactory. His first solution consisted of a long list of instructions based on what a human artist would do in certain situations. But this did not always work, partly because inevitably the list was open-ended.

Harold Cohen and his apprentice AARON making art.

Eventually, he found a way to teach AARON to use colors with a simple algorithm. We have limited ability to imagine differing chromatic arrangements, but our feedback system is terrific. A human artist can look at a picture as it evolves and decide exactly, say, what shade of yellow to add to a picture of sunflowers. AARON does not have a visual system at all, but Cohen devised a formula by which it can balance such factors as hue and saturation in any given image.

Can a machine ever be as creative as a Rembrandt or Picasso? To do that, Cohen argues, a robot would have to develop a sense of self. That may or may not happen, and "if it doesn't, it means that machines will never be creative in the same sense that humans are creative." The processes of such an artist involve an interplay between social, emotional, historical, psychological, and physiological factors that are dauntingly difficult to analyze, let alone replicate. This is what can give an image made by such an artist a deep level of meaning to a human eye.

One day, Cohen suggests, a machine might develop an equivalent sensibility, but even if that never comes to pass, "it doesn't mean that machines have no part to play with respect to creativity." As his own experience shows, artificial intelligence offers the artist something beyond an assistant or pupil: a new creative collaborator.

A new, expanded version of A Bigger Message, *Martin Gayford's book of conversations with David Hockney, will be published in May. His last story for* MIT Technology Review *was "Motion Pictures" (September/October 2015).*

HANK MORGAN/GETTY

TWELVE TOMORROWS

The 2016 edition of MIT Technology Review's science fiction anthology: visionary stories of the near future inspired by today's new technologies.

technologyreview.com/twelvetomorrows

MIT Technology Review

When Biology Meets Ideology

A new book reopens a notorious case of bungled science in the Soviet Union.

By Maggie Koerth-Baker

In 2012, the most important monastery in the Russian Orthodox Church published a biology textbook for 10th and 11th graders. It's called *General Biology*, but it's an explicitly creationist text, describing God's role in the natural world to counteract generations of official atheism in Russian schools. Darwinism, according to this book, has been disastrous for the world and for the Russian people in particular. It has led to an embrace of materialism, in both the philosophical and consumerist senses of the word. It's

CRISTIANA COUCEIRO

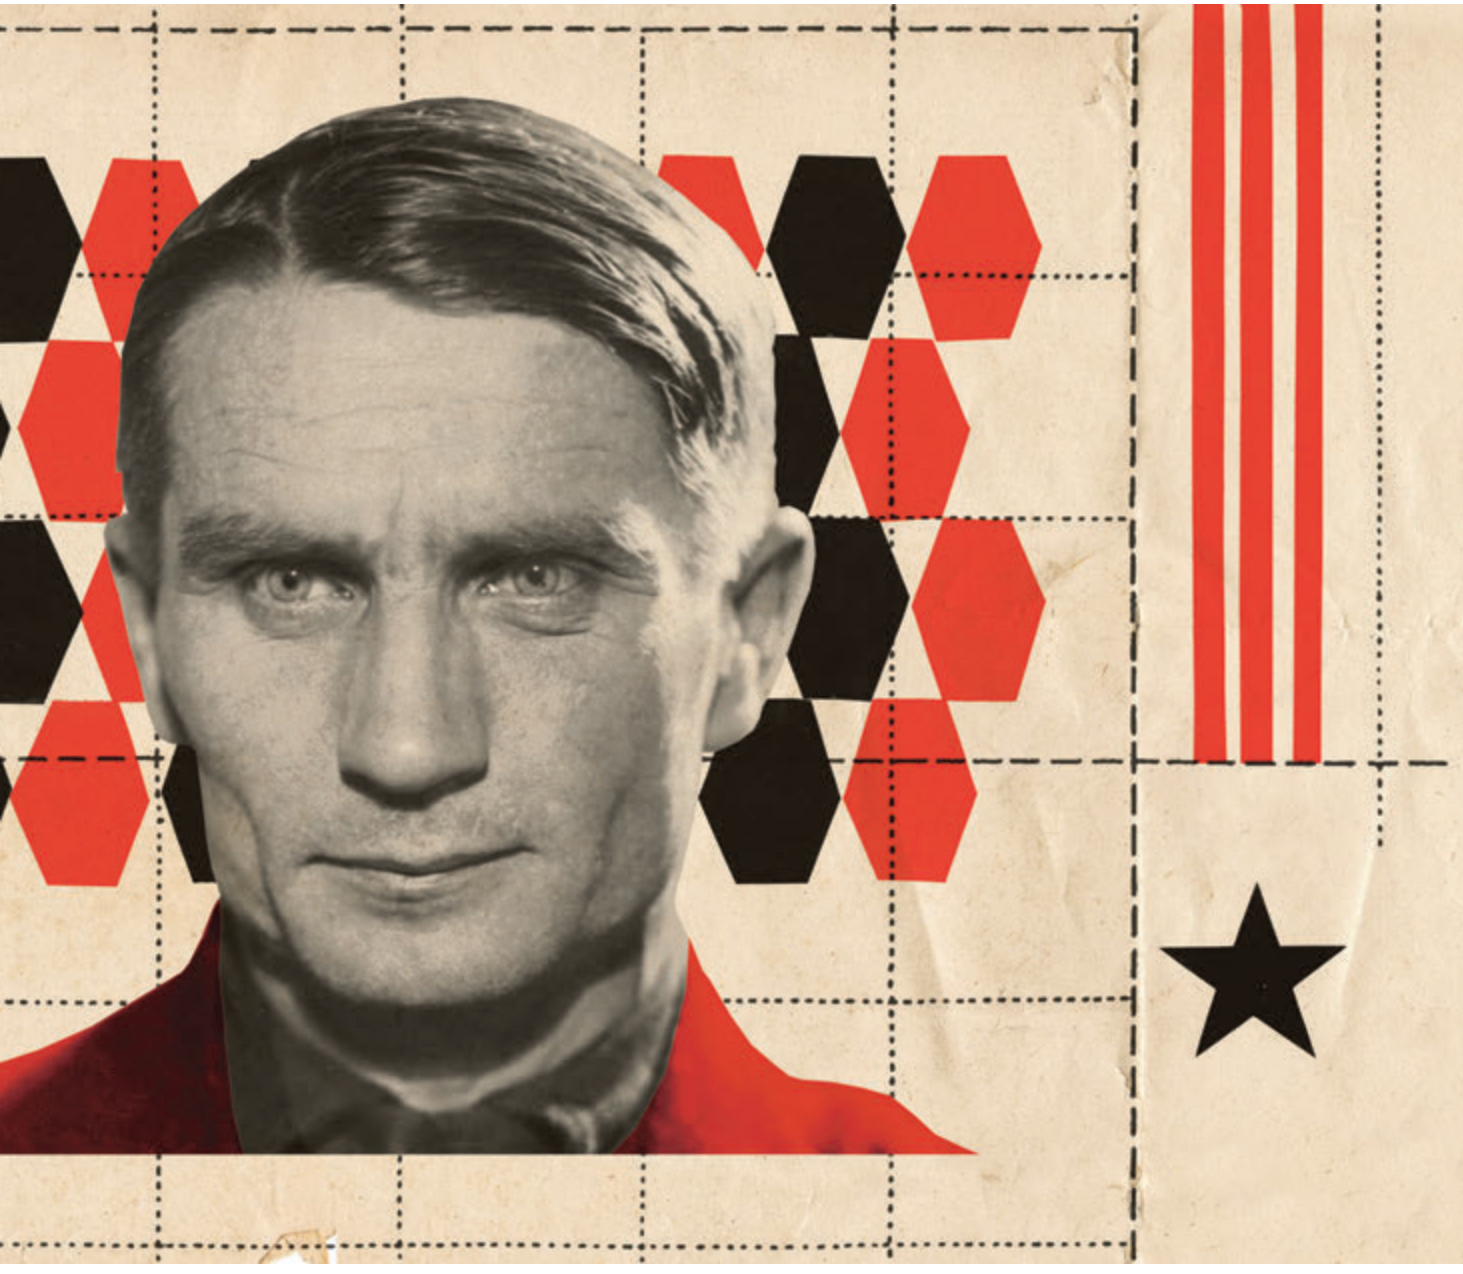

antithetical to Russian values because it's inherently intertwined with the dog-eat-dog lifestyles of 19th-century British capitalists. As the book denigrates natural selection, it praises the idea that characteristics acquired in one's life can be passed on to future generations. It refers to recent research on epigenetics, the study of how the environment affects genes' function in ways that are sometimes heritable.

Loren Graham, an MIT historian who has studied Russian science for decades, says *General Biology* is indicative of a recent resurgence of support for ideas once expounded by Trofim Lysenko, a Soviet biologist who rejected conventional genetics and tried to use acquired characteristics to improve agriculture. Lysenko set back farming and genetic research in the Soviet Union for decades, so why would anyone try to rehabilitate his ideas? Politics, essentially. In his new book, *Lysenko's Ghost*, Graham says *General Biology* is a reminder of "the continuing strength of the belief in the superiority of collectivism over individualism" in Russia.

That startled me. When I went to a fundamentalist Baptist high school in central Kansas, my ninth-grade biology textbook was, effectively, the American Protestant equivalent of *General Biology*. It, too, talked about acquired characteristics, but not as an alternative to Darwinism. Instead, we were taught that this theory, linked in the text to the French biologist Jean-Baptiste Lamarck, was inherently silly and served as proof

against evolution. Of course a giraffe that has to stretch to reach its food wouldn't produce babies with longer necks. Nor would a dog whose tail is docked have tail-less pups. It was just one more example of the ridiculous things evolutionists believed—beliefs that could be deeply dangerous. How dangerous? Well, everybody knows that Darwinism led people to reject God, abandon individual responsibility, and take up the mantle of collectivist communism in Russia.

Lysenko's Ghost: Epigenetics and Russia
By Loren Graham
Harvard University Press, 2016

The ways that politics, religion, cultural norms, and ideologies of all kinds distort science is at the heart of *Lysenko's Ghost*. Those ideologies can alter our interpretation of facts and reshape our understanding of natural events. They have the power to change the meanings of words, even scientific terms. All those issues are at the forefront as Graham explores whether modern epigenetic research—which indicates that environmental conditions like famine can affect gene expression and influence the health of people generations removed from the actual event—means that Lysenko's approach to agriculture was on the right track after all.

Spoiler: Lysenko has not been vindicated. Although epigenetics is deepening our understanding of how DNA works, it is not overturning the basic principles of genetic heredity that Lysenko challenged. Nonetheless, what to make of Lysenko now is a complicated question. As Graham points out, "the inheritance of acquired characteristics" did not mean the same thing to Lysenko—steeped in the politics and ethics of a collectivist Soviet Union—that it meant to Lamarck in France in the 1800s. It meant a third thing to many of Lysenko's Soviet science contemporaries, and something else entirely to the farmers and folk agronomists who thought they saw evidence of it long before Lamarck came along. Likewise, the name Lysenko means different things to Russians, Americans, and Europeans. "Natural selection" does not mean to modern biologists what it meant to the eugenicists of the 1930s. Even the word "true," Graham writes, is "thick and multidimensional." Graham calls this the contradiction between usage and accuracy.

Trofim Lysenko is a fascinating character. He was born a peasant in 1898. He rose to immense power in the 1940s

under Joseph Stalin by promoting a number of erroneous scientific techniques he claimed could increase wheat yields on famine-wracked collective farms. Among other things, he professed that by keeping seeds of winter wheat at low temperatures for longer than usual, he could convert the strain to a variety that would mature in the spring. When other scientists objected to his work, he attacked them in ways Graham calls "lethal and passive-aggressive," pointing them out to the secret police and letting the wheels of Stalinist "justice" do the rest. Not until the 1960s did he finally become a pariah, after the death of Stalin and the ouster of Nikita Khrushchev gave Lysenko's scientific foes an opportunity to denounce him as a fraud. Today, Lysenko is simultaneously a rallying point for a certain authoritarian strain of Russian nationalism and an embarrassment who leads Russian academics to avoid legitimate research on epigenetics.

Attempts to rehabilitate Lysenko's reputation remind us of how politics, religion, cultural norms, and ideologies of all kinds alter our interpretations of science.

Why was Lysenko opposed to the idea of inheritance through genes—and how did that mesh with Soviet ideology? Graham gives a partial answer. Even before Lysenko, in the 1920s, the German biologist Paul Kammerer and a slew of less-familiar Russian biologists promoted the idea of acquired characteristics as a sort of Marxist eugenics. In the West at this time, eugenics was all about creating a better society by making sure the "right" people (well-off and white) had lots of children and the "wrong" people (poor, disabled, black, and brown) had few or none. Kammerer, in contrast, promoted a eugenics based on improving environments. Marxism could make a better society by providing a better life, which would change the people who lived it, which would change their offspring. Over time, you would end up with an evolved human—the new "Soviet man," brighter and smarter and healthier than anything produced by simply pairing off generations of bourgeois capitalists.

The problem, of course, is that biology doesn't seem to play along. But Graham's narrative of how far Lysenko took these ideas is confusing. Lysenko did not actually believe that inheritance of acquired characteristics occurred in humans. And in Graham's telling, he seems to have been wishy-washy even on its applicability to agriculture.

That said, Graham is able to tell the story with intimate details. There's one particularly memorable anecdote in which a young Graham spots the aging, out-of-favor Lysenko at a posh Moscow restaurant in 1971 and maneuvers next to him at a shared table. Over a bowl of borscht, Graham introduces himself. He's uncomfortable, but he's certain he'll never get another crack at this.

Turns out Lysenko already knows who Graham is and doesn't like him. He feels Graham has unfairly fingered him as culpable in the deaths of many Russian biologists. In a remarkable back-and-forth, Graham and Lysenko argue over whether or not Lysenko was part of the oppressive Soviet system. They have no quibble about the facts. It's the meanings of the facts that they disagree about.

Maggie Koerth-Baker is a journalist and author in Minneapolis. Her work has appeared in Nature, Popular Mechanics, Gizmodo, *and the* New York Times.

29 Years Ago

How Technology Makes Us Obnoxious

A writer fretted that gadgets were leading to self-centered and rude behavior—decades before the smartphone.

"Rudeness is on the rise in the United States, and consumer technology is partly to blame. That may seem a harsh indictment for such seemingly innocuous creations as the telephone answering machine, boom box radio-cassette player, and talking computer chip in automobiles. But they are relatively cheap to manufacture and insinuate themselves into every corner of daily life. The result is a decline of civilization as we have known it.

As machines multiply our capacity to perform useful tasks, they boost our aptitude for self-centered action. Civilized behavior is predicated on the principle of one human being interacting with another, not a human being interacting with a mechanical or electronic extension of another person.

The simple telephone answering machine can turn into a devilish instrument if misused. Call screening can become the electronic equivalent of avoiding your neighbor's salutation on the street: the machine's owner can use it to avoid talking to anyone except the chosen few.

Yet it is hard to tell which enables people to be more rude: the answering machine that screens out obnoxious calls, or the telephone that permits intrusions at all times of day and night. A.G. Bell probably never envisioned aluminum-siding salespeople, college alumni associations, or lovelorn friends when he devised the telephone.

The Baby Bells are marketing a service that can make the answering machine seem like an instrument of civility. 'Call waiting' is handy for the small business, but in certain hands it permits the ultimate breach of telephone etiquette. A (former) friend used to chat happily with me on the phone until a second call came in from a potential suitor. After a few long waits, I would hang up and call back—only to reach her answering machine.

Personal-computer technology has spawned a vigorous rudeness that would be impossible without word processing. Once upon a time people wrote real letters to each other. Now every family with a home computer can generate its own 'personalized' holiday form letter. Practitioners of these 'arts' will defend them, no doubt, on the grounds that some communication is better than none. But is a pretense of personal rapport good enough?"

Excerpted from "Hey You! Make Way for My Technology!" by freelance writer David Lyon, in the August/September 1987 issue of Technology Review.

MIT Technology Review (ISSN 1099-274X), March/April 2016 Issue, Reg. U.S. Patent Office, is published bimonthly by MIT Technology Review, 1 Main St. Suite 13, Cambridge, MA 02142-1517. Periodicals postage paid at Boston, MA, and additional mailing offices. Postmaster: send address changes to MIT Technology Review, Subscriber Services, PO Box 5001, Big Sandy, TX 75755, or via the Internet at www.technologyreview.com/customerservice. Basic subscription rates: $41.94 per year within the United States; in all other countries, US$54. Publication Mail Agreement Number 40621028. Send undeliverable Canadian copies to PO Box 1051 Fort Erie, ON L2A 6C7. Printed in U.S.A. Audited by the Alliance for Audited Media

Aw, thanks **@bamadesigner**, we like you too.

We like to think Slack's changing the way that teams communicate. But don't take our word for it.
slack.com/love

slack
Work on purpose


TECHNOLOGY UNPLUGGED.
7QBT460
With Regen on Demand,™ location-based charging and four drive modes that let you choose how to use the battery power and backup gas-powered generator, the new Chevrolet Volt makes a pure electric drive possible!
THE NEW
CHEVROLET VOLT
FIND NEW ROADS™
1 EPA-estimated 53-mile EV range based on 106 MPGe combined city/highway (electric); 367-mile extended range based on 42 MPG combined city/highway (gas). Actual range varies with conditions.
CHEVROLET